

ETSI EN 302 878-4 V1.1.1 (2011-11)



**Access, Terminals, Transmission and Multiplexing (ATTM);
Third Generation Transmission Systems for
Interactive Cable Television Services - IP Cable Modems;
Part 4: MAC and Upper Layer Protocols;
DOCSIS 3.0**

Reference

DEN/ATTM-003006-4

Keywords

access, broadband, cable, data, IP, IPCable,
modem

ETSI

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

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

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

Important notice

Individual copies of the present document can be downloaded from:

<http://www.etsi.org>

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

Users of the present document should be aware that the document may be subject to revision or change of status.

Information on the current status of this and other ETSI documents is available at

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

If you find errors in the present document, please send your comment to one of the following services:

http://portal.etsi.org/chaicor/ETSI_support.asp

Copyright Notification

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

© European Telecommunications Standards Institute 2011.
All rights reserved.

DECTTM, PLUGTESTSTM, UMTSTM and the ETSI logo are Trade Marks of ETSI registered for the benefit of its Members.
3GPPTM and LTETM are Trade Marks of ETSI registered for the benefit of its Members and
of the 3GPP Organizational Partners.
GSM[®] and the GSM logo are Trade Marks registered and owned by the GSM Association.

Contents

Intellectual Property Rights	22
Foreword.....	22
1 Scope	23
1.1 Introduction and Purpose.....	23
1.2 Requirements.....	23
1.3 Conventions.....	23
2 References	23
2.1 Normative references	24
2.2 Informative references.....	27
3 Definitions and abbreviations.....	27
3.1 Definitions	27
3.2 Abbreviations	37
4 Void.....	41
5 Overview and Theory of Operations	41
5.1 DOCSIS 3.0 MULPI Key Features	41
5.2 Technical Overview	42
5.2.1 CMTS and CM Models.....	42
5.2.1.1 CMTS Model	42
5.2.1.1.1 Types of CMTS	42
5.2.1.1.2 CMTS Internal Forwarding Model.....	44
5.2.1.1.3 CMTS MAC Domain	45
5.2.1.2 CM Model	46
5.2.2 DOCSIS MAC Operation	46
5.2.2.1 QoS	46
5.2.2.1.1 Individual and Group Service Flows	47
5.2.2.2 Channel Bonding.....	48
5.2.2.2.1 Downstream Channel Bonding.....	48
5.2.2.2.2 Upstream Channel Bonding.....	49
5.2.2.3 Autonomous Load Balancing.....	51
5.2.3 Multicast Operation	51
5.2.4 Network and Higher Layer Protocols	52
5.2.5 CM and CPE Provisioning and Management	52
5.2.5.1 Initialization, Provisioning and Management of CMs	52
5.2.5.2 Initialization, Provisioning and Management of CPEs	53
5.2.6 Relationship to the Physical HFC Plant Topology	53
5.2.6.1 RF Topology Configuration.....	53
5.2.6.2 Frequency Assignment.....	55
5.2.7 Cable Service Group (CM-SG)	56
5.2.7.1 MAC Domain Channel Assignment.....	57
5.2.7.2 Multiple MAC Domains per Fibre Node	58
5.2.7.3 MAC Domain Downstream and Upstream Service Groups.....	59
5.2.7.4 Channel Bonding Topology Considerations	59
5.2.8 CMTS Downstream Service Model Example.....	60
6 Media Access Control Specification	61
6.1 Introduction	61
6.1.1 Overview	61
6.1.2 Definitions	61
6.1.2.1 MAC-Sublayer Domain	61
6.1.2.2 MAC Service Access Point.....	62
6.1.2.3 Service Flows	62
6.1.2.4 Upstream Intervals, Mini-Slots and 6,25-Microsecond Increments	62
6.1.2.4.1 TDMA mode	62
6.1.2.4.2 S-CDMA mode.....	63

6.1.2.5	MAC Frame	63
6.1.2.6	Logical Upstream Channels	63
6.1.2.6.1	Type 3 Logical Upstreams.....	64
6.1.2.6.2	Type 4 Logical Upstreams.....	64
6.1.3	Future Use.....	65
6.2	MAC Frame Formats.....	65
6.2.1	Generic MAC Frame Format	65
6.2.1.1	PMD Overhead	65
6.2.1.2	MAC Frame Transport.....	66
6.2.1.3	Ordering of Bits and Octets.....	66
6.2.1.3.1	Representing Negative Numbers	67
6.2.1.3.2	Type-Length-Value Fields.....	67
6.2.1.4	MAC Header Format.....	67
6.2.1.5	Data PDU	68
6.2.2	Packet-Based MAC Frames	69
6.2.2.1	Packet PDU and Isolation Packet PDU.....	69
6.2.3	ATM Cell MAC Frames	70
6.2.4	MAC-Specific Headers	70
6.2.4.1	Timing Header	70
6.2.4.2	MAC Management Header	71
6.2.4.3	Request Frame.....	71
6.2.4.4	Fragmentation Header.....	72
6.2.4.5	Queue-depth Based Request Frame	73
6.2.4.6	Concatenation Header	73
6.2.5	Extended MAC Headers	74
6.2.5.1	Piggyback Requests	76
6.2.5.2	Request Extended Header	76
6.2.5.3	Fragmentation Extended Header.....	76
6.2.5.4	Service Flow Extended Header.....	77
6.2.5.4.1	Payload Header Suppression Header.....	77
6.2.5.4.2	Unsolicited Grant Synchronization Header.....	78
6.2.5.5	BP_UP2 Extended Header.....	78
6.2.5.6	Downstream Service Extended Header.....	78
6.2.5.7	DPV Extended Header	80
6.3	Segment Header Format	80
6.4	MAC Management Messages.....	81
6.4.1	MAC Management Message Header	81
6.4.2	Time Synchronization (SYNC).....	84
6.4.3	Upstream Channel Descriptor (UCD).....	84
6.4.3.1	Example of UCD Encoded TLV Data.....	92
6.4.4	Upstream Bandwidth Allocation Map (MAP)	93
6.4.5	Ranging Request Messages	95
6.4.5.1	Ranging Request (RNG-REQ)	97
6.4.5.2	Initial Ranging Request (INIT-RNG-REQ)	98
6.4.5.3	Bonded Initial Ranging Request (B-INIT-RNG-REQ).....	99
6.4.5.3.1	Capability Flags.....	99
6.4.6	Ranging Response (RNG-RSP)	100
6.4.6.1	Encodings.....	101
6.4.6.2	Example of TLV Data.....	102
6.4.6.3	Transmit Equalization Encodings	103
6.4.6.4	RNG-RSP Channel Overrides.....	103
6.4.6.5	Upstream Channel Adjustments.....	104
6.4.6.6	T4 Timeout Multiplier.....	104
6.4.7	Registration Request Messages.....	105
6.4.7.1	Registration Request (REG-REQ)	106
6.4.7.2	Multipart Registration Request (REG-REQ-MP)	106
6.4.8	Registration Response Messages	107
6.4.8.1	Registration Response (REG-RSP).....	109
6.4.8.2	Multipart Registration Response (REG-RSP-MP).....	110
6.4.8.3	Encodings.....	111
6.4.8.3.1	Modem Capabilities.....	111
6.4.8.3.2	DOCSIS 1.0 Service Class Data.....	111

6.4.9	Registration Acknowledge (REG-ACK)	112
6.4.10	Upstream Channel Change Request (UCC-REQ).....	113
6.4.11	Upstream Channel Change Response (UCC-RSP).....	114
6.4.12	Dynamic Service Addition - Request (DSA-REQ).....	114
6.4.12.1	CM-Initiated Dynamic Service Addition	115
6.4.12.2	CMTS-Initiated Dynamic Service Addition.....	115
6.4.13	Dynamic Service Addition - Response (DSA-RSP)	116
6.4.13.1	CM-Initiated Dynamic Service Addition	117
6.4.13.2	CMTS-Initiated Dynamic Service Addition.....	117
6.4.14	Dynamic Service Addition - Acknowledge (DSA-ACK)	117
6.4.15	Dynamic Service Change - Request (DSC-REQ).....	118
6.4.16	Dynamic Service Change - Response (DSC-RSP).....	120
6.4.17	Dynamic Service Change - Acknowledge (DSC-ACK)	121
6.4.18	Dynamic Service Deletion - Request (DSD-REQ).....	122
6.4.19	Dynamic Service Deletion - Response (DSD-RSP).....	122
6.4.20	Dynamic Channel Change - Request (DCC-REQ)	123
6.4.20.1	Encodings.....	124
6.4.20.1.1	Upstream Channel ID	124
6.4.20.1.2	Downstream Parameters	124
6.4.20.1.3	Initialization Technique	125
6.4.20.1.4	UCD Substitution	126
6.4.20.1.5	Security Association Identifier (SAID) Substitution	127
6.4.20.1.6	Service Flow Substitutions	127
6.4.20.1.7	CMTS MAC Address	128
6.4.21	Dynamic Channel Change - Response (DCC-RSP).....	129
6.4.21.1	Encodings.....	129
6.4.21.1.1	CM Jump Time.....	130
6.4.22	Dynamic Channel Change - Acknowledge (DCC-ACK)	130
6.4.23	Device Class Identification Request (DCI-REQ)	131
6.4.24	Device Class Identification Response (DCI-RSP).....	132
6.4.25	Upstream Transmitter Disable (UP-DIS).....	132
6.4.26	Test Request (TST-REQ).....	134
6.4.27	Downstream Channel Descriptor (DCD).....	135
6.4.28	MAC Domain Descriptor (MDD).....	135
6.4.28.1	TLV Encodings	136
6.4.28.1.1	Downstream Active Channel List TLV	136
6.4.28.1.2	MAC Domain Downstream Service Group (MD-DS-SG) TLV	137
6.4.28.1.3	Downstream Ambiguity Resolution Frequency List TLV.....	138
6.4.28.1.4	Receive Channel Profile Reporting Control TLV	138
6.4.28.1.5	IP Initialization Parameters TLV	139
6.4.28.1.6	Early Authentication and Encryption (EAE) Enable/Disable TLV	139
6.4.28.1.7	Upstream Active Channel List TLV	140
6.4.28.1.8	Upstream Ambiguity Resolution Channel List TLV	140
6.4.28.1.9	Upstream Frequency Range TLV	141
6.4.28.1.10	Symbol Clock Locking Indicator	141
6.4.28.1.11	CM-STATUS Event Control	141
6.4.28.1.12	Upstream Transmit Power Reporting	142
6.4.28.1.13	DSG DA-to-DSID Association Entry	142
6.4.28.1.14	CM-STATUS Event Enable for Non-Channel-Specific Events	142
6.4.28.1.15	Extended Upstream Transmit Power Support	143
6.4.29	Dynamic Bonding Change Request (DBC-REQ)	143
6.4.30	Dynamic Bonding Change Response (DBC-RSP)	144
6.4.31	Dynamic Bonding Change Acknowledge (DBC-ACK)	146
6.4.32	DOCSIS Path Verify Request (DPV-REQ)	146
6.4.33	DOCSIS Path Verify Response (DPV-RSP)	148
6.4.34	Status Report (CM-STATUS).....	148
6.4.34.1	TLV Encodings	149
6.4.35	CM Control Request (CM-CTRL-REQ).....	149
6.4.35.1	TLV Encodings	150
6.4.36	CM Control Response (CM-CTRL-RSP).....	150
7	Media Access Control Protocol Operation.....	151

7.1	Timing and Synchronization	151
7.1.1	Global Timing Reference.....	151
7.1.2	CM Synchronization	152
7.1.3	Ranging.....	152
7.1.3.1	Broadcast Initial Ranging.....	152
7.1.3.2	Unicast Initial Ranging	153
7.1.4	Timing Units and Relationships.....	153
7.1.4.1	TDMA Timing Units and Relationships	153
7.1.4.1.1	Mini-Slot Capacity	153
7.1.4.1.2	Mini-Slot Numbering	154
7.1.4.2	S-CDMA Timing Units and Relationships	154
7.1.4.2.1	Mini-Slot Capacity	154
7.1.4.2.2	Mini-Slot Numbering	154
7.2	Upstream Data Transmission	155
7.2.1	Upstream Bandwidth Allocation.....	155
7.2.1.1	The Allocation MAP MAC Management Message	156
7.2.1.2	Information Elements.....	156
7.2.1.2.1	The Request IE	156
7.2.1.2.2	The Request/Data IE.....	157
7.2.1.2.3	The Initial Maintenance IE	157
7.2.1.2.4	The Station Maintenance IE	157
7.2.1.2.5	Short and Long Data Grant IEs	157
7.2.1.2.6	Data Acknowledge IE.....	158
7.2.1.2.7	Expansion IE	158
7.2.1.2.8	Null IE	158
7.2.1.2.9	Advanced PHY Short and Long Data Grant IEs	158
7.2.1.2.10	Advanced PHY Unsolicited Grant IE.....	159
7.2.1.3	Requesting with Multiple Transmit Channel Mode Disabled	159
7.2.1.4	Requesting with Multiple Transmit Channel Mode Enabled	160
7.2.1.4.1	Request Mechanisms for Segment Header OFF Service Flows	160
7.2.1.4.2	Request Mechanisms for Segment Header ON Service Flows.....	160
7.2.1.5	Information Element Feature Usage Summary	165
7.2.1.6	Map Transmission and Timing	165
7.2.1.7	Protocol Example.....	166
7.2.1.8	MAP Generation Example - Two Logical Upstreams.....	167
7.2.2	Upstream Transmission and Contention Resolution.....	168
7.2.2.1	Contention Resolution Overview	168
7.2.2.1.1	Contention Resolution with Multiple Transmit Channel Mode Disabled	168
7.2.2.1.2	Contention Resolution with Multiple Transmit Channel Mode Enabled.....	169
7.2.2.2	Transmit Opportunities	171
7.2.2.3	CM Bandwidth Utilization.....	172
7.2.3	Upstream Service Flow Scheduling Services	172
7.2.3.1	Unsolicited Grant Service	173
7.2.3.2	Real-Time Polling Service	173
7.2.3.3	Unsolicited Grant Service with Activity Detection.....	174
7.2.3.4	Non-Real-Time Polling Service	175
7.2.3.5	Best Effort Service	175
7.2.3.6	Other Services	175
7.2.3.6.1	Committed Information Rate (CIR).....	175
7.2.3.7	Parameter Applicability for Upstream Service Scheduling.....	175
7.2.3.8	CM Transmit Behavior	176
7.2.4	Continuous Concatenation and Fragmentation	176
7.2.5	Pre-3.0 DOCSIS Concatenation and Fragmentation.....	177
7.2.5.1	Concatenation.....	177
7.2.5.2	Fragmentation	178
7.2.5.2.1	CM Fragmentation Support.....	180
7.2.5.2.2	CMTS Fragmentation Support	182
7.2.5.2.3	Fragmentation Example.....	183
7.2.5.2.4	Pre-Registration Fragmentation.....	185
7.2.5.2.5	Considerations for Concatenated Packets and Fragmentation.....	186
7.3	Upstream - Downstream Channel Association within a MAC Domain	186
7.3.1	Primary Downstream Channels	186

7.3.2	MAP and UCD Messages	187
7.3.3	Multiple MAC Domains	187
7.4	DSID Definition	187
7.5	Quality of Service	188
7.5.1	Concepts	188
7.5.1.1	Service Flows	188
7.5.1.2	Classifiers	190
7.5.1.2.1	Upstream and Downstream QoS Classifiers	190
7.5.1.2.2	Upstream Drop Classifiers	192
7.5.2	Object Model	192
7.5.3	Service Classes	194
7.5.4	Authorization	195
7.5.5	States of Service Flows	195
7.5.5.1	Deferred Service Flows	195
7.5.5.1.1	Provisioned Service Flows	196
7.5.5.1.2	Authorized Service Flows	196
7.5.5.2	Admitted Service Flows	196
7.5.5.3	Active Service Flows	197
7.5.6	Service Flows and Classifiers	197
7.5.6.1	Policy-Based Classification and Service Classes	198
7.5.7	General Operation	198
7.5.7.1	Static Operation	198
7.5.7.2	Dynamic Service Flow Creation - CM Initiated	199
7.5.7.3	Dynamic Service Flow Creation - CMTS Initiated	200
7.5.7.4	Dynamic Service Flow Modification and Deletion	200
7.5.8	QoS Support for Joined IP Multicast Traffic	201
7.5.8.1	Overview	202
7.5.8.2	Group Configuration and Group QoS Configuration Tables	203
7.5.8.3	Instantiating Group Classifier Rules and Group Service Flows	204
7.5.8.3.1	Examples of GCR and GSF Instantiation	206
7.5.8.4	Default Group Service Flows	212
7.5.8.5	Service Class QoS Parameter Changes	212
7.5.8.6	Group QoS Configuration Changes	212
7.5.9	Other Multicast and Broadcast Traffic	213
7.6	Downstream Traffic Priority	213
7.6.1	Traffic Priority Ordering and Mapping to CM Output Queues	213
7.7	Payload Header Suppression	214
7.7.1	Overview	214
7.7.1.1	PHSI-indexed PHS	215
7.7.1.2	DSID-indexed PHS	215
7.7.2	Example Applications	216
7.7.3	Operation	216
7.7.4	Signalling	218
7.7.4.1	Signalling PHSI-Indexed Payload Header Suppression	218
7.7.4.2	Signalling DSID-Indexed Payload Header Suppression	219
7.7.5	Payload Header Suppression Examples	220
7.7.5.1	Upstream Example	220
7.7.5.2	Downstream Example	221
7.7.5.3	DSID-Indexed Multicast Example	222
7.8	Data Link Encryption Support	223
7.8.1	MAC Messages	223
7.8.2	Framing	223
7.8.3	Multiple Transmit Channel Mode Operation and Packet Encryption	223
8	Channel Bonding	223
8.1	Upstream and Downstream Common Aspects	224
8.1.1	Service Flow Assignment	224
8.1.2	CMTS Bonding and Topology Requirements	227
8.2	Downstream Channel Bonding	228
8.2.1	Multiple Downstream Channel Overview	228
8.2.2	CMTS Downstream Bonding Operation	229
8.2.3	Sequenced Downstream Packets	229

8.2.3.1	Downstream Sequencing.....	230
8.2.3.2	Skew Requirements.....	232
8.2.3.3	Resequencing DSID Signalling.....	233
8.2.4	Cable Modem Physical Receive Channel Configuration.....	233
8.2.4.1	Receive Channels.....	234
8.2.4.2	Receive Modules.....	234
8.2.4.2.1	Receive Module Interconnection.....	236
8.2.4.3	Receive Channel Profile.....	237
8.2.4.3.1	Standard Receive Channel Profiles	237
8.2.4.4	Receive Channel Configuration	238
8.2.4.4.1	Static Receive Module Assignments	239
8.2.4.5	RCC Message Sequence Example	240
8.2.5	QoS for Downstream Channel Bonding	241
8.3	Upstream Channel Bonding	241
8.3.1	Granting Bandwidth.....	241
8.3.2	Upstream Transmissions with Upstream Channel Bonding	241
8.3.2.1	Segment Header ON Operation.....	241
8.3.2.2	Segment Header OFF Operation	242
8.3.3	Dynamic Range Window.....	242
8.3.3.1	Channels Added During Registration	242
8.3.3.2	Channels Added by a DBC-REQ	243
8.3.3.3	Channels Deleted by a DBC-REQ	244
8.3.3.4	UCD Changes Burst Profiles Resulting in New Value for P_{hi}	244
8.3.3.5	Power Offset in RNG-RSP Causing Dynamic Range Window Violation	244
8.4	Partial Service	244
9	Data Forwarding.....	245
9.1	General Forwarding Requirements.....	245
9.1.1	CMTS Forwarding Rules.....	246
9.1.1.1	General CMTS Forwarding.....	246
9.1.1.2	DSID Labeling	247
9.1.2	CM Address Acquisition, Filtering and Forwarding Rules.....	247
9.1.2.1	MAC Address Acquisition.....	248
9.1.2.2	CM Filtering Rules.....	248
9.1.2.3	CM Forwarding Rules.....	249
9.1.2.3.1	CM Pre-Operational Forwarding Behavior	249
9.1.2.3.2	CM Operational Forwarding Behavior.....	249
9.2	Multicast Forwarding	251
9.2.1	Introduction.....	251
9.2.2	Downstream Multicast Forwarding	252
9.2.2.1	Examples of Downstream Multicast Forwarding using DSIDs	253
9.2.2.2	Labeling Multicast Packets with DSIDs	255
9.2.2.2.1	Mixed CM environment	256
9.2.2.2.2	Pre-Registration DSID.....	256
9.2.2.2.3	Upstream Multicast Traffic from a Multicast Client	256
9.2.2.3	Communicating DSIDs and group forwarding attributes to a CM.....	256
9.2.2.4	DSID based Filtering and Forwarding by a Cable Modem	257
9.2.2.5	Individually Directed Multicast	258
9.2.3	Downstream Multicast Traffic Encryption	258
9.2.3.1	Multicast Encryption Overview	258
9.2.3.2	Dynamic Multicast Encryption	259
9.2.3.3	DSIDs and SAIDs	259
9.2.3.4	Pre-Registration Multicast Encryption.....	259
9.2.4	Static Multicast Session Encodings	260
9.2.5	IGMP and MLD Support	260
9.2.5.1	Motivation behind taking CM out of IGMP Control Plane.....	260
9.2.5.2	IP multicast service model support	260
9.2.5.3	IGMP and MLD membership handling.....	261
9.2.5.4	IGMPv2/MLDv1 Leave Processing.....	262
9.2.5.5	IGMP and MLD version and query support.....	262
9.2.5.6	Separation of Query Domains.....	262
9.2.6	Encrypted Multicast Downstream Forwarding Example	263

9.2.7	IP Multicast Join Authorization	266
9.2.7.1	Maximum Multicast Sessions	266
9.2.7.2	Session Rules	267
9.2.7.2.1	IP Multicast Profiles	267
9.2.7.2.2	Static IP Multicast Join Authorization Rules.....	268
9.2.7.3	CM Configuration File	268
9.2.7.3.1	IP Multicast Profile Name Subtype	268
9.2.7.3.2	Static IP Multicast Session Rule Subtype.....	268
9.2.7.4	Matching Session Rules	268
9.2.7.5	IP Multicast Profile Changes.....	269
10	Cable Modem - CMTS Interaction.....	269
10.1	CMTS Initialization.....	269
10.2	Cable Modem Initialization and Reinitialization.....	270
10.2.1	Scan for Downstream Channel	270
10.2.2	Continue Downstream Scanning.....	271
10.2.3	Service Group Discovery and Initial Ranging	271
10.2.3.1	Read MAC Domain Descriptor (MDD).....	273
10.2.3.2	MDDs Not Found on Primary Downstream.....	274
10.2.3.3	Determination of MD-DS-SG	275
10.2.3.4	Ranging Holdoff	277
10.2.3.5	Determination of MD-US-SG	279
10.2.3.5.1	Bonded Initial Ranging.....	280
10.2.3.5.2	Continue US Ambiguity Initial Ranging	282
10.2.3.6	Obtain Upstream Parameters / Try Next Upstream (DOCSIS 2.0 Initialization).....	284
10.2.3.6.1	Message Flows During Scanning and Upstream Parameter Acquisition.....	285
10.2.3.7	Ranging and Automatic Adjustments.....	286
10.2.3.7.1	Adjust Transmit Parameters	290
10.2.3.8	CMTS Determination of Cable Modem Service Group and Initial Ranging	291
10.2.4	Authentication.....	293
10.2.5	Establishing IP Connectivity.....	293
10.2.5.1	Establish IPv4 Network Connectivity	301
10.2.5.1.1	DHCPv4 Fields Used by the CM	302
10.2.5.1.2	Use of T1 and T2 Timers.....	303
10.2.5.1.3	CMTS Requirements	303
10.2.5.2	Establish IPv6 Network Connectivity	304
10.2.5.2.1	Obtain Link-Local Address	305
10.2.5.2.2	Obtain default routers	305
10.2.5.2.3	Obtain IPv6 management address and other configuration parameters.....	305
10.2.5.2.4	Use of T1 and T2 Timers.....	306
10.2.5.2.5	CMTS Requirements	307
10.2.5.3	Alternate Provisioning Mode (APM) Operation	307
10.2.5.4	Dual-stack Provisioning Mode (DPM).....	308
10.2.5.5	Establish Time of Day.....	308
10.2.5.6	Transfer Operational Parameters.....	309
10.2.5.7	Configuration File Processing.....	310
10.2.5.8	Post-registration Failures to Renew IP Addresses.....	311
10.2.6	Registration with the CMTS	311
10.2.6.1	Cable Modem Requirements	311
10.2.6.2	CMTS Requirements.....	319
10.2.6.2.1	Channel Assignment During Registration	325
10.2.7	Baseline Privacy Initialization	327
10.2.8	Service IDs During CM Initialization.....	327
10.3	Periodic Maintenance	328
10.4	Fault Detection and Recovery	330
10.4.1	CM Downstream Channel Interruptions	331
10.4.2	MAC Layer Error-Handling	332
10.4.2.1	Error Recovery During Pre-3.0 DOCSIS Fragmentation.....	333
10.4.2.2	Error Recovery During Segmentation with Segment Headers On	333
10.4.3	CM Status Report	334
10.4.3.1	Event Codes	337
10.5	DOCSIS Path Verification	340

10.5.1	DPV Overview.....	340
10.5.2	DPV Reference Points	340
10.5.3	DPV Math.....	342
10.5.4	DPV Per Path Operation.....	342
10.5.4.1	DPV Ping	343
10.5.5	DPV Per Packet Operation	343
11	Dynamic Operations.....	344
11.1	Upstream Channel Descriptor Changes.....	344
11.2	Dynamic Service Flow Changes	345
11.2.1	Dynamic Service Flow State Transitions.....	346
11.2.2	Dynamic Service Addition.....	354
11.2.2.1	CM Initiated Dynamic Service Addition.....	354
11.2.2.2	CMTS Initiated Dynamic Service Addition.....	355
11.2.2.3	Dynamic Service Addition State Transition Diagrams	356
11.2.3	Dynamic Service Change.....	364
11.2.3.1	CM-Initiated Dynamic Service Change	365
11.2.3.2	CMTS-Initiated Dynamic Service Change.....	365
11.2.3.3	Dynamic Service Change State Transition Diagrams	367
11.2.4	Dynamic Service Deletion	375
11.2.4.1	CM Initiated Dynamic Service Deletion	375
11.2.4.2	CMTS Initiated Dynamic Service Deletion	375
11.2.4.3	Dynamic Service Deletion State Transition Diagrams.....	376
11.3	Pre-3.0 DOCSIS Upstream Channel Changes.....	380
11.4	Dynamic Downstream and/or Upstream Channel Changes	383
11.4.1	DCC General Operation.....	383
11.4.1.1	Derivation of T15 Timer	385
11.4.1.2	Initialization Technique	386
11.4.1.2.1	Initialization Technique Zero (0).....	386
11.4.1.2.2	Initialization Technique One (1).....	387
11.4.1.2.3	Initialization Technique Two (2).....	387
11.4.1.2.4	Initialization Technique Three (3).....	387
11.4.1.2.5	Initialization Technique Four (4).....	387
11.4.2	DCC Exception Conditions	388
11.4.3	DCC State Transition Diagrams	389
11.4.4	DCC Performance.....	396
11.5	Dynamic Bonding Change (DBC).....	397
11.5.1	DBC General Operation.....	397
11.5.1.1	Changes to the Receive Channel Set.....	398
11.5.1.2	Changes to a DSID.....	399
11.5.1.2.1	Changes to Resequencing Encodings	399
11.5.1.2.2	Changes to Multicast Encodings	400
11.5.1.3	Changes to the Security Association for encrypting downstream traffic	402
11.5.1.4	Changes to the Transmit Channel Set	402
11.5.1.4.1	Impact of TCS Changes on Periodic Ranging	403
11.5.1.4.2	Exception Conditions for TCS Changes.....	403
11.5.1.5	Changes to the Service Flow SID Cluster Assignments.....	404
11.5.1.5.1	Bandwidth Sufficiency	405
11.5.1.6	Initialization Technique	405
11.5.1.6.1	Initialization Technique One (1).....	405
11.5.1.6.2	Initialization Technique Two (2).....	406
11.5.1.6.3	Initialization Technique Three (3).....	406
11.5.1.6.4	Initialization Technique Four (4).....	406
11.5.1.7	Fragmentation of DBC-REQ Messages	407
11.5.2	Exception Conditions.....	407
11.5.3	DBC State Transition Diagrams	408
11.5.3.1	CMTS DBC State Transition Diagrams.....	408
11.5.3.2	CM DBC State Transition Diagrams	413
11.6	Autonomous Load Balancing.....	418
11.6.1	Load Balancing Groups	419
11.6.1.1	General Load Balancing Groups	419
11.6.1.2	Restricted Load Balancing Groups.....	420

11.6.2	CMTS Load Balancing Operation	420
11.6.3	Multiple Channel Load Balancing	421
11.6.4	Initialization Techniques.....	421
11.6.5	Load Balancing Policies	421
11.6.6	Load Balancing Priorities	422
11.6.7	Load Balancing and Multicast	422
11.6.8	Externally-Directed Load Balancing	423
12	Supporting Future New Cable Modem Capabilities	423
12.1	Downloading Cable Modem Operating Software	423
12.2	Future Capabilities	424
Annex A (normative):	Well_known_Addresses.....	425
A.1	Addresses	425
A.1.1	General MAC Addresses.....	425
A.1.2	Well-known IPv6 Addresses	425
A.2	MAC Service IDs	425
A.2.1	All CMs and No CM Service IDs.....	425
A.2.2	Well-Known Multicast Service IDs	426
A.2.3	Priority Request Service IDs	426
A.3	MPEG PID	426
Annex B (normative):	Parameters and Constants	427
Annex C (normative):	Common TLV Encodings.....	430
C.1	Encodings for Configuration and MAC-Layer Messaging	431
C.1.1	Configuration File and Registration Settings	431
C.1.1.1	Downstream Frequency Configuration Setting.....	431
C.1.1.2	Upstream Channel ID Configuration Setting.....	432
C.1.1.3	Network Access Control Object	432
C.1.1.4	DOCSIS 1.0 Class of Service Configuration Setting.....	432
C.1.1.4.1	Class ID.....	433
C.1.1.4.2	Maximum Downstream Rate Configuration Setting	433
C.1.1.4.3	Maximum Upstream Rate Configuration Setting.....	433
C.1.1.4.4	Upstream Channel Priority Configuration Setting	434
C.1.1.4.5	Guaranteed Minimum Upstream Channel Data Rate Configuration Setting	434
C.1.1.4.6	Maximum Upstream Channel Transmit Burst Configuration Setting	434
C.1.1.4.7	Class-of-Service Privacy Enable	434
C.1.1.5	CM Message Integrity Check (MIC) Configuration Setting.....	435
C.1.1.6	CMTS Message Integrity Check (MIC) Configuration Setting	435
C.1.1.7	Maximum Number of CPEs	435
C.1.1.8	TFTP Server Timestamp.....	435
C.1.1.9	TFTP Server Provisioned Modem IPv4 Address.....	435
C.1.1.10	TFTP Server Provisioned Modem IPv6 Address.....	436
C.1.1.11	Upstream Packet Classification Configuration Setting	436
C.1.1.12	Downstream Packet Classification Configuration Setting	436
C.1.1.13	Upstream Service Flow Encodings	436
C.1.1.14	Downstream Service Flow Encodings	436
C.1.1.15	Payload Header Suppression.....	436
C.1.1.16	Maximum Number of Classifiers.....	436
C.1.1.17	Privacy Enable	437
C.1.1.18	DOCSIS Extension Field	437
C.1.1.18.1	General Extension Information	437
C.1.1.18.1.1	CM Load Balancing Policy ID	438
C.1.1.18.1.2	CM Load Balancing Priority	438
C.1.1.18.1.3	CM Load Balancing Group ID	438
C.1.1.18.1.4	CM Ranging Class ID Extension.....	438
C.1.1.18.1.5	L2VPN Encoding	438
C.1.1.18.1.6	Extended CMTS MIC Configuration Setting	439
C.1.1.18.1.7	Source Address Verification (SAV) Authorization Encoding.....	440

C.1.1.18.1.8	Cable Modem Attribute Masks.....	441
C.1.1.18.1.9	IP Multicast Join Authorization Encoding	442
C.1.1.18.1.10	Service Type Identifier	444
C.1.1.18.2	Vendor-Specific Information	444
C.1.1.19	Subscriber Management TLVs	445
C.1.1.19.1	Subscriber Management Control.....	445
C.1.1.19.2	Subscriber Management CPE IPv4 List.....	445
C.1.1.19.3	Subscriber Management CPE IPv6 Prefix List	445
C.1.1.19.4	Subscriber Management Filter Groups.....	445
C.1.1.19.5	Subscriber Management Control Max CPE IPv6 Addresses	446
C.1.1.19.6	Subscriber Management CPE IPv6 List.....	446
C.1.1.20	Enable 2.0 Mode.....	446
C.1.1.21	Enable Test Modes.....	446
C.1.1.22	Downstream Channel List	447
C.1.1.22.1	Single Downstream Channel	447
C.1.1.22.1.1	Single Downstream Channel Timeout	448
C.1.1.22.1.2	Single Downstream Channel Frequency	448
C.1.1.22.2	Downstream Frequency Range	448
C.1.1.22.2.1	Downstream Frequency Range Timeout	448
C.1.1.22.2.2	Downstream Frequency Range Start	448
C.1.1.22.2.3	Downstream Frequency Range End	449
C.1.1.22.2.4	Downstream Frequency Range Step Size.....	449
C.1.1.22.3	Default Scanning.....	449
C.1.1.22.4	Examples Illustrating Usage of the Downstream Channel List.....	449
C.1.1.23	Static Multicast MAC Address	450
C.1.1.24	Downstream Unencrypted Traffic (DUT) Filtering Encoding.....	450
C.1.1.25	Channel Assignment Configuration Settings.....	451
C.1.1.25.1	Transmit Channel Assignment Configuration Setting.....	451
C.1.1.25.2	Receive Channel Assignment Configuration Setting	451
C.1.1.26	Upstream Drop Classifier Group ID	451
C.1.1.27	CMTS Static Multicast Session Encoding	451
C.1.1.27.1	Static Multicast Group Encoding	452
C.1.1.27.2	Static Multicast Source Encoding	452
C.1.1.27.3	Static Multicast CMIM Encoding	452
C.1.2	Configuration-File-Specific Settings.....	452
C.1.2.1	End-of-Data Marker.....	452
C.1.2.2	Pad Configuration Setting.....	452
C.1.2.3	Software Upgrade Filename	453
C.1.2.4	SNMP Write-Access Control.....	453
C.1.2.5	SNMP MIB Object	453
C.1.2.6	CPE Ethernet MAC Address	454
C.1.2.7	Software Upgrade IPv4 TFTP Server	454
C.1.2.8	Software Upgrade IPv6 TFTP Server	454
C.1.2.9	SnmpV3 Kickstart Value	454
C.1.2.9.1	SnmpV3 Kickstart Security Name	454
C.1.2.9.2	SnmpV3 Kickstart Manager Public Number.....	455
C.1.2.10	Manufacturer Code Verification Certificate	455
C.1.2.11	Co-signer Code Verification Certificate	455
C.1.2.12	SNMPv3 Notification Receiver	455
C.1.2.12.1	SNMPv3 Notification Receiver IPv4 Address	455
C.1.2.12.2	SNMPv3 Notification Receiver UDP Port Number.....	456
C.1.2.12.3	SNMPv3 Notification Receiver Trap Type.....	456
C.1.2.12.4	SNMPv3 Notification Receiver Timeout	456
C.1.2.12.5	SNMPv3 Notification Receiver Retries	456
C.1.2.12.6	SNMPv3 Notification Receiver Filtering Parameters	456
C.1.2.12.7	SNMPv3 Notification Receiver Security Name.....	456
C.1.2.12.8	SNMPv3 Notification Receiver IPv6 Address	457
C.1.2.13	SNMPv1v2c Coexistence Configuration.....	457
C.1.2.13.1	SNMPv1v2c Community Name.....	457
C.1.2.13.2	SNMPv1v2c Transport Address Access	457
C.1.2.13.2.1	SNMPv1v2c Transport Address.....	457
C.1.2.13.2.2	SNMPv1v2c Transport Address Mask	458

C.1.2.13.3	SNMPv1v2c Access View Type	458
C.1.2.13.4	SNMPv1v2c Access View Name	458
C.1.2.14	SNMPv3 Access View Configuration	458
C.1.2.14.1	SNMPv3 Access View Name	459
C.1.2.14.2	SNMPv3 Access View Subtree	459
C.1.2.14.3	SNMPv3 Access View Mask	459
C.1.2.14.4	SNMPv3 Access View Type	459
C.1.2.15	SNMP CPE Access Control	459
C.1.2.16	Management Event Control Encoding	460
C.1.3	Registration-Request/Response-Specific Encodings	460
C.1.3.1	Modem Capabilities Encoding	460
C.1.3.1.1	Concatenation Support	460
C.1.3.1.2	DOCSIS Version	460
C.1.3.1.3	Fragmentation Support	461
C.1.3.1.4	Payload Header Suppression Support	461
C.1.3.1.5	IGMP Support	461
C.1.3.1.6	Privacy Support	461
C.1.3.1.7	Downstream SAID Support	461
C.1.3.1.8	Upstream Service Flow Support	461
C.1.3.1.9	Optional Filtering Support	462
C.1.3.1.10	Transmit Pre-Equalizer Taps per Modulation Interval	462
C.1.3.1.11	Number of Transmit Equalizer Taps	462
C.1.3.1.12	DCC Support	462
C.1.3.1.13	IP Filters Support	463
C.1.3.1.14	LLC Filters Support	463
C.1.3.1.15	Expanded Unicast SID Space	463
C.1.3.1.16	Ranging Hold-Off Support	463
C.1.3.1.17	L2VPN Capability	464
C.1.3.1.18	L2VPN eSAFE Host Capability	464
C.1.3.1.19	Downstream Unencrypted Traffic (DUT) Filtering	464
C.1.3.1.20	Upstream Frequency Range Support	464
C.1.3.1.21	Upstream Symbol Rate Support	464
C.1.3.1.22	Selectable Active Code Mode 2 Support	465
C.1.3.1.23	Code Hopping Mode 2 Support	465
C.1.3.1.24	Multiple Transmit Channel Support	465
C.1.3.1.25	5,12 Msps Upstream Transmit Channel Support	466
C.1.3.1.26	2.56 Msps Upstream Transmit Channel Support	466
C.1.3.1.27	Total SID Cluster Support	466
C.1.3.1.28	SID Clusters per Service Flow Support	466
C.1.3.1.29	Multiple Receive Channel Support	466
C.1.3.1.30	Total Downstream Service ID (DSID) Support	467
C.1.3.1.31	Resequencing Downstream Service ID (DSID) Support	467
C.1.3.1.32	Multicast Downstream Service ID (DSID) Support	467
C.1.3.1.33	Multicast DSID Forwarding	467
C.1.3.1.34	Frame Control Type Forwarding Capability	468
C.1.3.1.35	DPV Capability	468
C.1.3.1.36	Unsolicited Grant Service/Upstream Service Flow Support	468
C.1.3.1.37	MAP and UCD Receipt Support	468
C.1.3.1.38	Upstream Drop Classifier Support	469
C.1.3.1.39	IPv6 Support	469
C.1.3.1.40	Extended Upstream Transmit Power Capability	469
C.1.3.2	Vendor ID Encoding	470
C.1.3.3	Modem IP Address	470
C.1.3.4	Service(s) Not Available Response	470
C.1.3.5	Vendor-Specific Capabilities	470
C.1.3.6	CM Initialization Reason	471
C.1.4	Dynamic-Message-Specific Encodings	471
C.1.4.1	HMAC-Digest	471
C.1.4.2	Authorization Block	472
C.1.4.3	Key Sequence Number	472
C.1.5	Registration, Dynamic Service and Dynamic Bonding Settings	472
C.1.5.1	Transmit Channel Configuration (TCC)	472

C.1.5.1.1	Transmit Channel Configuration (TCC) Reference	473
C.1.5.1.2	Upstream Channel Action	473
C.1.5.1.3	Upstream Channel ID	474
C.1.5.1.4	New Upstream Channel ID	474
C.1.5.1.5	UCD	474
C.1.5.1.6	Ranging SID	474
C.1.5.1.7	Initialization Technique	475
C.1.5.1.8	Ranging Parameters	475
C.1.5.1.8.1	Ranging Reference Channel ID	475
C.1.5.1.8.2	Timing Offset, Integer Part	476
C.1.5.1.8.3	Timing Offset, Fractional Part	476
C.1.5.1.9	Dynamic Range Window	476
C.1.5.1.10	TCC Error Encodings	477
C.1.5.1.10.1	Reported Parameter	477
C.1.5.1.10.2	Error Code	477
C.1.5.1.10.3	Error Message	477
C.1.5.2	Service Flow SID Cluster Assignments	477
C.1.5.2.1	SFID	478
C.1.5.2.2	SID Cluster Encoding	478
C.1.5.2.3	SID Cluster Switchover Criteria	479
C.1.5.2.3.1	Maximum Requests per SID Cluster	479
C.1.5.2.3.2	Maximum Outstanding Bytes per SID Cluster	479
C.1.5.2.3.3	Maximum Total Bytes Requested per SID Cluster	479
C.1.5.3	CM Receive Channel (RCP/RCC) Encodings	480
C.1.5.3.1	RCP-ID	480
C.1.5.3.2	RCP Name	481
C.1.5.3.3	RCP Center Frequency Spacing	481
C.1.5.3.4	Receive Module Encoding	481
C.1.5.3.4.1	Receive Module Index	481
C.1.5.3.4.2	Receive Module Adjacent Channels	482
C.1.5.3.4.3	Receive Module Channel Block Range	482
C.1.5.3.4.4	Receive Module First Channel Center Frequency Assignment	482
C.1.5.3.4.5	Receive Module Resequencing Channel Subset Capability	483
C.1.5.3.4.6	Receive Module Connectivity	483
C.1.5.3.4.7	Receive Module Common Physical Layer Parameter	483
C.1.5.3.5	Receive Channels	483
C.1.5.3.5.1	Receive Channel Index	483
C.1.5.3.5.2	Receive Channel Connectivity	484
C.1.5.3.5.3	Receive Channel Connected Offset	484
C.1.5.3.5.4	Receive Channel Center Frequency Assignment	484
C.1.5.3.5.5	Receive Channel Primary Downstream Channel Indicator	485
C.1.5.3.6	Partial Service Downstream Channels	485
C.1.5.3.7	Receive Channel Profile/Configuration Vendor Specific Parameters	485
C.1.5.3.8	RCC Error Encodings	485
C.1.5.3.8.1	Receive Module or Receive Channel	485
C.1.5.3.8.2	Receive Module Index or Receive Channel Index	486
C.1.5.3.8.3	Reported Parameter	486
C.1.5.3.8.4	Error Code	486
C.1.5.3.8.5	Error Message	486
C.1.5.4	DSID Encodings	486
C.1.5.4.1	Downstream Service Identifier (DSID)	486
C.1.5.4.2	Downstream Service Identifier Action	487
C.1.5.4.3	Downstream Resequencing Encodings	487
C.1.5.4.3.1	Resequencing DSID	487
C.1.5.4.3.2	Downstream Resequencing Channel List	487
C.1.5.4.3.3	DSID Resequencing Wait Time	487
C.1.5.4.3.4	Resequencing Warning Threshold	488
C.1.5.4.3.5	CM-STATUS Maximum Event Hold-Off Timer for Sequence Out-of-Range Events	488
C.1.5.4.4	Multicast Encodings	488
C.1.5.4.4.1	Client MAC Address Encodings	488
C.1.5.4.4.2	Multicast CM Interface Mask	489
C.1.5.4.4.3	Multicast Group MAC Addresses Encodings	489

C.1.5.4.4.4	Payload Header Suppression Encodings	489
C.1.5.5	Security Association Encoding	490
C.1.5.5.1	SA Action.....	490
C.1.5.5.2	SA-Descriptor	490
C.1.5.6	Initializing Channel Timeout	491
C.2	Quality-of-Service-Related Encodings.....	491
C.2.1	Packet Classification Encodings.....	491
C.2.1.1	Upstream Packet Classification Encoding	491
C.2.1.2	Upstream Drop Packet Classification Encoding	491
C.2.1.3	Downstream Packet Classification Encoding	491
C.2.1.4	General Packet Classifier Encodings	492
C.2.1.4.1	Classifier Reference	492
C.2.1.4.2	Classifier Identifier.....	492
C.2.1.4.3	Service Flow Reference	492
C.2.1.4.4	Service Flow Identifier.....	492
C.2.1.4.5	Rule Priority	492
C.2.1.4.6	Classifier Activation State.....	493
C.2.1.4.7	Dynamic Service Change Action	493
C.2.1.4.8	CM Interface Mask (CMIM) Encoding.....	493
C.2.1.5	Classifier Error Encodings	493
C.2.1.5.1	Errored Parameter	494
C.2.1.5.2	Error Code.....	494
C.2.1.5.3	Error Message	494
C.2.1.6	IPv4 Packet Classification Encodings.....	495
C.2.1.6.1	IPv4 Type of Service Range and Mask	495
C.2.1.6.2	IP Protocol.....	495
C.2.1.6.3	IPv4 Source Address	495
C.2.1.6.4	IPv4 Source Mask	495
C.2.1.6.5	IPv4 Destination Address.....	496
C.2.1.6.6	IPv4 Destination Mask.....	496
C.2.1.7	TCP/UDP Packet Classification Encodings.....	496
C.2.1.7.1	TCP/UDP Source Port Start	496
C.2.1.7.2	TCP/UDP Source Port End	496
C.2.1.7.3	TCP/UDP Destination Port Start.....	496
C.2.1.7.4	TCP/UDP Destination Port End	497
C.2.1.8	Ethernet LLC Packet Classification Encodings	497
C.2.1.8.1	Destination MAC Address	497
C.2.1.8.2	Source MAC Address	497
C.2.1.8.3	Ethertype/DSAP/MacType	497
C.2.1.9	IEEE 802.1P/Q Packet Classification Encodings	498
C.2.1.9.1	IEEE 802.1P User_Priority	498
C.2.1.9.2	IEEE 802.1Q VLAN_ID	498
C.2.1.10	IPv6 Packet Classification Encodings.....	499
C.2.1.10.1	IPv6 Traffic Class Range and Mask.....	499
C.2.1.10.2	IPv6 Flow Label.....	499
C.2.1.10.3	IPv6 Next Header Type.....	499
C.2.1.10.4	IPv6 Source Address	500
C.2.1.10.5	IPv6 Source Prefix Length (bits).....	500
C.2.1.10.6	IPv6 Destination Address.....	500
C.2.1.10.7	IPv6 Destination Prefix Length (bits)	500
C.2.1.11	Vendor Specific Classifier Parameters	501
C.2.2	Service Flow Encodings	501
C.2.2.1	Upstream Service Flow Encodings.....	501
C.2.2.2	Downstream Service Flow Encodings	501
C.2.2.3	General Service Flow Encodings	501
C.2.2.3.1	Service Flow Reference	501
C.2.2.3.2	Service Flow Identifier.....	502
C.2.2.3.3	Service Identifier	502
C.2.2.3.4	Service Class Name.....	502
C.2.2.3.5	Quality of Service Parameter Set Type	502
C.2.2.3.6	Service Flow Required Attribute Mask.....	503

C.2.2.3.7	Service Flow Forbidden Attribute Mask	503
C.2.2.3.8	Service Flow Attribute Aggregation Rule Mask	504
C.2.2.3.9	Application Identifier	504
C.2.2.4	Service Flow Error Encodings	504
C.2.2.4.1	Errored Parameter	505
C.2.2.4.2	Error Code	505
C.2.2.4.3	Error Message	505
C.2.2.5	Common Upstream and Downstream Quality-of-Service Parameter Encodings	505
C.2.2.5.1	Traffic Priority	505
C.2.2.5.2	Maximum Sustained Traffic Rate	506
C.2.2.5.2.1	Upstream Maximum Sustained Traffic Rate	506
C.2.2.5.2.2	Downstream Maximum Sustained Traffic Rate	506
C.2.2.5.3	Maximum Traffic Burst	507
C.2.2.5.4	Minimum Reserved Traffic Rate	507
C.2.2.5.5	Assumed Minimum Reserved Rate Packet Size	508
C.2.2.5.6	Timeout for Active QoS Parameters	508
C.2.2.5.7	Timeout for Admitted QoS Parameters	508
C.2.2.5.8	Vendor Specific QoS Parameters	509
C.2.2.5.9	IP Type Of Service (DSCP) Overwrite	509
C.2.2.5.10	Peak Traffic Rate	509
C.2.2.5.10.1	Upstream Peak Traffic Rate	509
C.2.2.5.10.2	Downstream Peak Traffic Rate	510
C.2.2.6	Upstream-Specific QoS Parameter Encodings	510
C.2.2.6.1	Maximum Concatenated Burst	510
C.2.2.6.2	Service Flow Scheduling Type	511
C.2.2.6.3	Request/Transmission Policy	511
C.2.2.6.4	Nominal Polling Interval	512
C.2.2.6.5	Tolerated Poll Jitter	512
C.2.2.6.6	Unsolicited Grant Size	513
C.2.2.6.7	Nominal Grant Interval	513
C.2.2.6.8	Tolerated Grant Jitter	513
C.2.2.6.9	Grants per Interval	514
C.2.2.6.10	Unsolicited Grant Time Reference	514
C.2.2.6.11	Multiplier to Contention Request Backoff Window	514
C.2.2.6.12	Multiplier to Number of Bytes Requested	514
C.2.2.7	Downstream-Specific QoS Parameter Encodings	515
C.2.2.7.1	Maximum Downstream Latency	515
C.2.2.7.2	Downstream Resequencing	515
C.2.2.8	Payload Header Suppression	515
C.2.2.8.1	Classifier Reference	515
C.2.2.8.2	Classifier Identifier	516
C.2.2.8.3	Service Flow Reference	516
C.2.2.8.4	Service Flow Identifier	516
C.2.2.8.5	Dynamic Service Change Action	516
C.2.2.8.6	Dynamic Bonding Change Action	516
C.2.2.9	Payload Header Suppression Error Encodings	517
C.2.2.9.1	Errored Parameter	517
C.2.2.9.2	Error Code	517
C.2.2.9.3	Error Message	518
C.2.2.10	Payload Header Suppression Rule Encodings	518
C.2.2.10.1	Payload Header Suppression Field (PHSF)	518
C.2.2.10.2	Payload Header Suppression Index (PHSI)	518
C.2.2.10.3	Payload Header Suppression Mask (PHSM)	519
C.2.2.10.4	Payload Header Suppression Size (PHSS)	519
C.2.2.10.5	Payload Header Suppression Verification (PHSV)	519
C.2.2.10.6	Vendor Specific PHS Parameters	520
C.3	Encodings for Other Interfaces	520
C.3.1	Baseline Privacy Configuration Settings Option	520
C.4	Confirmation Code	520

Annex D (normative):	CM Configuration Interface Specification	525
D.1	CM Configuration	525
D.1.1	CM Binary Configuration File Format.....	525
D.1.2	Configuration File Settings.....	525
D.1.3	Configuration File Creation.....	527
D.1.3.1	CM MIC Calculation	528
D.2	Configuration Verification	529
D.2.1	CMTS MIC Calculation	529
D.2.1.1	Pre-3.0 DOCSIS CMTS MIC Digest Calculation.....	530
D.2.1.2	Extended CMTS MIC Digest Calculation	530
Annex E (normative):	Standard Receive Channel Profile Encodings.....	532
Annex F (normative):	The DOCSIS MAC/PHY Interface (DMPI)	538
F.1	Scope	538
F.2	Conventions.....	538
F.2.1	Terminology	538
F.2.2	Ordering of Bits and Bytes	538
F.2.3	Signal Naming Conventions.....	538
F.2.4	Active Clock Edge.....	538
F.2.5	Timing Specifications.....	539
F.3	Overview	539
F.3.1	Downstream Data	540
F.3.2	Upstream Data	541
F.3.3	Upstream Control	541
F.3.4	SPI Bus.....	541
F.4	Signals	542
F.4.1	Downstream Data	542
F.5	Upstream Data.....	542
F.5.1	Upstream Control	543
F.5.2	SPI Bus.....	543
F.5.3	Parity	543
F.5.3.1	Downstream Data	543
F.5.3.2	Upstream Data	544
F.5.3.3	Upstream Control.....	544
F.5.4	Interrupts	544
F.6	Protocol	544
F.6.1	Downstream Data (ITU-T Recommendation J.83, annex A)	544
F.6.2	Downstream Data (ITU-T Recommendation J.83, annex B)	545
F.6.3	Upstream Data	545
F.6.4	Upstream Control	546
F.6.4.1	Counter Synchronization	546
F.6.4.2	Upstream Control Messages	547
F.6.5	SPI Bus.....	548
F.7	Electrical Specifications	548
F.7.1	DC Specifications.....	548
F.8	Timing Specifications.....	549
F.8.1	Downstream Data	549
F.8.2	Upstream Data	549
F.8.3	Upstream Control	549
F.8.4	SPI Bus.....	550
F.9	Data Format and Usage	550
F.9.1	Downstream Data	550
F.9.2	Upstream Data	550
F.9.2.1	Block Format	550

F.9.2.2	FIRST_DATA Block.....	551
F.9.2.3	MIDDLE_DATA Block.....	551
F.9.2.4	LAST_DATA Block.....	551
F.9.2.5	PHY_STATUS Block.....	552
F.9.2.6	NO_BURST Block.....	552
F.9.2.7	CHANNEL Block.....	553
F.9.2.8	Block Usage.....	553
F.9.2.8.1	Overview.....	553
F.9.2.8.2	Burst Data Transfer.....	553
F.9.2.8.3	No Burst Status Transfer.....	554
F.9.2.8.4	UCD Change Indication.....	554
F.9.2.8.5	Logical Channel Support.....	554
F.9.3	Upstream Control.....	555
F.9.3.1	Interval Description Message.....	555
F.9.3.2	UCD Change Message.....	556
F.9.4	SPI Bus.....	557

Annex G (normative): Compatibility with Previous Versions of DOCSIS558

G.1	General Interoperability Issues.....	558
G.1.1	Initial Ranging.....	558
G.1.2	Topology Resolution.....	558
G.1.3	Early Authentication and Encryption (EAE).....	559
G.1.4	Provisioning.....	559
G.1.5	Registration.....	563
G.1.6	Requesting Bandwidth.....	568
G.1.7	Encryption Support.....	568
G.1.8	Downstream Channel Bonding.....	568
G.1.9	Upstream Channel Bonding and Transmit Channel Configuration Support.....	568
G.1.10	Dynamic Service Establishment.....	569
G.1.11	Fragmentation.....	569
G.1.12	Multicast Support.....	569
G.1.13	Changing Upstream Channels.....	569
G.1.14	Changing Downstream Channels.....	570
G.2	Support for Hybrid Devices.....	570
G.3	Upstream Physical Layer Interoperability.....	571
G.3.1	DOCSIS 2.0 TDMA Interoperability.....	571
G.3.1.1	Mixed-mode operation with TDMA on a Type 2 channel.....	571
G.3.1.2	Interoperability and Performance.....	571
G.3.2	DOCSIS 2.0 S-CDMA Interoperability.....	572
G.3.2.1	Mixed mode operation with S-CDMA.....	572
G.3.2.2	Interoperability and Performance.....	572
G.3.3	DOCSIS 3.0 Interoperability.....	572
G.4	Multicast Support for Interaction with Pre-3.0 DOCSIS Devices.....	573
G.4.1	Multicast DSID Forwarding (MDF) Capability Exchange.....	573
G.4.2	GMAC-Explicit Multicast DSID Forwarding Mode.....	573
G.4.2.1	Example: Forwarding of Multicast Traffic to a Client behind a GMAC-Explicit CM.....	575
G.4.2.2	GMAC-Promiscuous Override.....	577
G.4.3	MDF Mode 0.....	577
G.4.3.1	CMTS Requirements with MDF Mode 0.....	577
G.4.3.2	CM Requirements with MDF Disabled.....	578
G.4.3.2.1	Requirements for IGMP Management.....	578
G.4.3.2.1.1	IGMP Timer Requirements.....	579
G.4.3.2.2	Forwarding Requirements for User Joined Multicast.....	579
G.4.3.2.3	Forwarding Requirements For Multicast Traffic Associated with IPv6.....	581

Annex H (normative): DHCPv6 Vendor Specific Information Options for DOCSIS 3.0.....582

Annex I: Void583

Annex J (normative):	DHCPv4 Vendor Identifying Vendor Specific Options for DOCSIS 3.0	584
Annex K (normative):	The Data-Over-Cable Spanning Tree Protocol.....	585
K.1	Background	585
K.2	Public Spanning Tree	585
K.3	Public Spanning Tree Protocol Details	586
K.4	Spanning Tree Parameters and Defaults.....	587
K.4.1	Path Cost	587
K.4.2	Bridge Priority	587
Annex L (informative):	MAC Service Definition	588
L.1	MAC Service Overview	588
L.1.1	MAC Service Parameters	589
L.2	MAC Data Service Interface	589
L.2.1	MAC_DATA_INDIVIDUAL.request.....	590
L.2.1.1	Databases	590
L.2.1.2	Pseudocode	591
L.2.2	MAC_DATA_GROUP.request.....	592
L.2.3	MAC_DATA_INTERNAL.request	593
L.2.4	MAC_GRANT_SYNCHRONIZE.indicate	593
L.2.5	MAC_CMTS_MASTER_CLOCK_SYNCHRONIZE.indicate.....	594
L.3	MAC Control Service Interface	594
L.3.1	MAC_REGISTRATION_RESPONSE.indicate	594
L.3.2	MAC_CREATE_SERVICE_FLOW.request	594
L.3.3	MAC_CREATE_SERVICE_FLOW.response.....	595
L.3.4	MAC_CREATE_SERVICE_FLOW.indicate	595
L.3.5	MAC_DELETE_SERVICE_FLOW.request.....	595
L.3.6	MAC_DELETE_SERVICE_FLOW.response	596
L.3.7	MAC_DELETE_SERVICE_FLOW.indicate	596
L.3.8	MAC_CHANGE_SERVICE_FLOW.request	596
L.3.9	MAC_CHANGE_SERVICE_FLOW.response.....	596
L.3.10	MAC_CHANGE_SERVICE_FLOW.indicate.....	596
L.4	MAC Service Usage Scenarios	597
L.4.1	Transmission of PDUs from Upper Layer Service to MAC DATA Service.....	597
L.4.2	Reception of PDUs to Upper Layer Service from MAC DATA Service	597
L.4.3	Sample Sequence of MAC Control and MAC Data Services	597
Annex M (informative):	Plant Topologies.....	599
M.1	Single Downstream and Single Upstream per Cable Segment.....	599
M.2	Multiple Downstreams and Multiple Upstreams per Cable Segment	601
M.2.1	HFC Plant Topologies	602
M.2.2	Normal Operation.....	603
M.2.3	Initial Ranging.....	603
M.2.4	Dynamic Channel Change.....	604
Annex N (informative):	DOCSIS Transmission and Contention Resolution.....	605
N.1	Multiple Transmit Channel Mode	605
N.1.1	Introduction	605
N.1.2	Variable Definitions	606
N.1.3	State Examples	607
N.1.3.1	Idle - Waiting for a Packet to Transmit.....	607
N.1.3.2	Grant Pending - Waiting for a Grant.....	607
N.1.3.3	Deferring - Determine Proper Transmission Timing and Transmit	607
N.1.4	Function Examples	607

N.1.4.1	CalcDefer() - Determine Defer Amount	607
N.1.4.2	UtilizeGrant() - Determine Best Use of a Grant	608
N.1.4.3	Retry()	609
N.1.4.4	Process Map()	609
N.1.4.5	timeout (sid)	609
N.1.4.6	is_my_SID(sid)	609
N.2	Non-Multiple Transmit Channel Mode	609
N.2.1	Introduction	609
N.2.2	Variable Definitions	610
N.2.3	State Examples	611
N.2.3.1	Idle - Waiting for a Packet to Transmit	611
N.2.3.2	Data Ack Pending - Waiting for Data Ack only	611
N.2.3.3	Grant Pending - Waiting for a Grant	611
N.2.3.4	Deferring - Determine Proper Transmission Timing and Transmit	611
N.2.4	Function Examples	612
N.2.4.1	CalcDefer() - Determine Defer Amount	612
N.2.4.2	UtilizeGrant() - Determine Best Use of a Grant	612
N.2.4.3	Retry()	613
Annex O (informative):	Unsolicited Grant Services	614
O.1	Unsolicited Grant Service (UGS)	614
O.1.1	Introduction	614
O.1.2	Configuration Parameters	614
O.1.3	Operation	614
O.1.4	Jitter	615
O.1.5	Synchronization Issues	615
O.2	Unsolicited Grant Service with Activity Detection (UGS-AD)	616
O.2.1	Introduction	616
O.2.2	MAC Configuration Parameters	616
O.2.3	Operation	616
O.2.4	Example	617
O.2.5	Talk Spurt Grant Burst	617
O.2.6	Admission Considerations	618
O.3	Multiple Transmit Channel Mode Considerations for Unsolicited Grant Services	619
Annex P (informative):	Error Recovery Examples	620
Annex Q (informative):	SDL Notation	622
Annex R (informative):	Notes on Address Configuration in DOCSIS 3.0	623
Annex S (informative):	IP Multicast Replication Examples	624
S.1	Scenario I: First Multicast Client joiner to a multicast session (Start of a new Multicast Session)	624
S.1.1	Scenario I - Case 1	625
S.1.2	Scenario I - Case 2	626
S.1.3	Scenario I - Case 3	627
S.2	Scenario II: A Multicast Client joining an existing multicast session that is being forwarded bonded, with FC-Type 10 (Typical 3.0 Multicast Mode of Operation)	628
S.2.1	Scenario II - Case 1	628
S.2.2	Scenario II - Case 2	631
S.2.3	Scenario II - Case 3	632
Annex T (informative):	IGMP Example for DOCSIS 2.0 Backwards Compatibility Mode	634
T.1	Events	634
T.2	Actions	635
Annex U (informative):	CM Multicast DSID Filtering Summary	636

Annex V (informative):	Example DHCPv6 Solicit Message Contents.....	637
Annex W (informative):	Dynamic Operations Examples	638
W.1	Dynamic Channel Change Example Operation.....	638
W.1.1	Example Signalling	638
W.1.2	Example Timing	640
W.1.2.1	Upstream and Downstream Change (Use Channel Directly: CMTS Supplies All TLV Hints).....	640
W.1.2.2	Upstream and Downstream Change (Station maintenance: CMTS Supplies No TLV Hints)	641
W.2	Dynamic Bonding Change Example Operation	642
W.2.1	Change to Transmit Channel Set and Service Flow SID Cluster Assignments	642
W.2.2	Change to Receive Channel Set and Downstream Resequencing Channel List.....	643
W.3	Autonomous Load Balancing Example.....	644
Annex X (informative):	Bibliography.....	647
History	648

Intellectual Property Rights

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

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are or may be or may become, essential to the present document.

Foreword

This European Standard (EN) has been produced by ETSI Technical Committee Access, Terminals, Transmission and Multiplexing (ATTM).

The present document is part 4 of a multi-part deliverable. Full details of the entire series can be found in part 1 [11].

National transposition dates	
Date of adoption of this EN:	14 November 2011
Date of latest announcement of this EN (doa):	29 February 2012
Date of latest publication of new National Standard or endorsement of this EN (dop/e):	31 August 2012
Date of withdrawal of any conflicting National Standard (dow):	31 August 2012

1 Scope

1.1 Introduction and Purpose

The present document is part of a series of specifications that define the third generation of high-speed data-over-cable systems. This series was developed for the benefit of the cable industry and includes contributions by operators and vendors from North America, Europe and other regions.

The present document defines the Media Access Control (MAC) layer protocols of DOCSIS 3.0 as well as requirements for upper layer protocols (e.g. IP, DHCP, etc.). DOCSIS 3.0 introduces new MAC layer features beyond what were present in earlier versions of DOCSIS.

1.2 Requirements

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

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

The present document defines many features and parameters and a valid range for each parameter is usually specified. Equipment (CM and CMTS) requirements are always explicitly stated. Equipment must comply with all mandatory (MUST and MUST NOT) requirements to be considered compliant with this specification. Support of non-mandatory features and parameter values is optional.

1.3 Conventions

In this specification the following convention applies any time a bit field is displayed in a figure. The bit field should be interpreted by reading the figure from left to right, then from top to bottom, with the MSB being the first bit so read and the LSB being the last bit so read.

MIB syntax and XML Schema syntax is represented by this code sample font.

2 References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the reference document (including any amendments) applies.

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

NOTE: While any hyperlinks included in this clause were valid at the time of publication ETSI cannot guarantee their long term validity.

2.1 Normative references

The following referenced documents are necessary for the application of the present document.

- [1] Cable Television Laboratories, Inc CL-SP-CANN-DHCP-Reg-I04-100611 (June 2010): "CableLabs® Specifications - CableLabs' DHCP Options Registry".
- [2] Cable Television Laboratories, Inc. CM-SP-DEPI-I08-100611 (June 2010): "Data-Over-Cable Service Interface Specifications - Modular Headend Architecture - Downstream External PHY Interface Specification".
- [3] Cable Television Laboratories, Inc. CM-SP-CMCIV3.0-I01-080320 (March 2008): "Data-Over-Cable Service Interface Specifications - Cable Modem to Customer Premise Equipment - Interface Specification".
- [4] ETSI EN 302 878-3: "Access, Terminals, Transmission and Multiplexing (ATTM); Third Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems; Part 3: Downstream Radio Frequency Interface; DOCSIS 3.0".
- [5] Cable Television Laboratories, Inc. CM-SP-DSG-I15-100611 (June 2010): "Data-Over-Cable Service Interface Specifications - DOCSIS Set-top Gateway (DSG) Interface Specification".
- [6] Cable Televisions Laboratories, Inc. CM-SP-DTI-I05-081209 (December 2008): "Data-Over-Cable Service Interface Specifications - Modular Headend Architecture - DOCSIS Timing Interface Specification".
- [7] Cable Television Laboratories, Inc. CM-SP-eDOCSIS-I20-100611 (June 2010): "Data-Over-Cable Service Interface Specifications - eDOCSIS™ Specification".
- [8] Cable Television Laboratories, Inc. CM-SP-L2VPN-I09-100611 (June 2010): "Data-Over-Cable Service Interface Specifications - Business Services over DOCSIS® - Layer 2 Virtual Private Networks".
- [9] Cable Television Laboratories, Inc. CM-SP-OSSIV2.0-C01-081104 (November 2008): "Data-Over-Cable Service Interface Specifications - DOCSIS 2.0 - Operations Support System Interface Specification".
- [10] Cable Television Laboratories, Inc. CM-SP-OSSIV3.0-I13-101008 (October 2010): "Data-Over-Cable Service Interface Specifications - DOCSIS 3.0 - Operations Support System Interface Specification".
- [11] ETSI EN 302 878-1: "Access, Terminals, Transmission and Multiplexing (ATTM); Third Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems; Part 1: General; DOCSIS 3.0".
- [12] ETSI EN 302 878-2: "Access, Terminals, Transmission and Multiplexing (ATTM); Third Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems; Part 2: Physical Layer; DOCSIS 3.0".
- [13] Cable Television Laboratories, Inc. CM-SP-RFIV2.0-C02-090422 (April 2009): "Data-Over-Cable Service Interface Specifications - DOCSIS 2.0 - Radio Frequency Interface Specification".
- [14] ETSI ES 202 488-2 (V1.1.1): "Access and Terminals (AT); Second Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems; Part 2: Radio frequency interface specification".
- [15] ETSI EN 302 878-5: "Access, Terminals, Transmission and Multiplexing (ATTM); Third Generation Transmission Systems for Interactive Cable Television Services - IP Cable Modems; Part 5: Security Services; DOCSIS 3.0".
- [16] IEEE 802.1D-2004: "IEEE Standard for Local and metropolitan area networks -- Media Access Control (MAC) Bridges" (Incorporates IEEE 802.1t-2001 and IEEE 802.1w).

- [17] ISO/IEC 8802-2:1998: "Information technology -- Telecommunications and information exchange between systems -- Local and metropolitan area networks -- Specific requirements -- Part 2: Logical link control".
- [18] ISO/IEC 8802-3:2000: "Information technology -- Telecommunications and information exchange between systems -- Local and metropolitan area networks -- Specific requirements -- Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications".
- [19] ISO/IEC 8825-1:2008: "Information technology -- ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)".
- [20] ISO/IEC 15802-3:1998: "Information technology -- Telecommunications and information exchange between systems -- Local and metropolitan area networks -- Common specifications -- Part 3: Media Access Control (MAC) Bridges".
- [21] ITU-T Recommendation X.25 (1998 - Corrigendum 1): "Interface between Data Terminal Equipment (DTE) and Data Circuit-terminating Equipment (DCE) for terminals operating in the packet mode and connected to public data networks by dedicated circuit".
- [22] ITU-T Recommendation Z.100 (2007): "Specification and Description Language (SDL)".
- [23] Cable Television Laboratories, Inc. PKT-SP-MM-I05-091029 (October 2009): "PacketCable™ Specification - Multimedia Specification".
- [24] Cable Television Laboratories, Inc. PKT-SP-DQoS-C01-071129 (November 2007): "PacketCable™ Dynamic Quality-of-Service Specification".
- [25] Cable Television Laboratories, Inc. PKT-SP-EC-MGCP-C01-071129 (November 2007): "PacketCable™ Network-Based Call Signalling - Protocol Specification".
- [26] IETF RFC 768 (1980): "User Datagram Protocol", J. Postel.
- [27] IETF RFC 868/STD0026 (1983): "Time Protocol".
- [28] IETF RFC 1042 (1988): "Standard for the transmission of IP datagrams over IEEE 802 networks".
- [29] IETF RFC 1112 (1989): "Host extensions for IP multicasting".
- [30] IETF RFC 1157 (1990): "Simple Network Management Protocol (SNMP)".
- [31] IETF RFC 1350 (1992): "The TFTP Protocol (Revision 2)".
- [32] IETF RFC 4188 (2005): "Definitions of Managed Objects for Bridges".
- [33] IETF RFC 3232 (2002): "Assigned Numbers: RFC 1700 is Replaced by an On-line Database".
- [34] IETF RFC 1812 (1995): "Requirements for IP Version 4 Routers".
- [35] IETF RFC 1945 (1996): "Hypertext Transfer Protocol -- HTTP/1.0".
- [36] IETF RFC 2104 (1997): "HMAC: Keyed-Hashing for Message Authentication".
- [37] IETF RFC 2131 (1997): "Dynamic Host Configuration Protocol".
- [38] IETF RFC 2132 (1997): "DHCP Options and BOOTP Vendor Extensions".
- [39] IETF RFC 2236 (1997): "Internet Group Management Protocol, Version 2".
- [40] IETF RFC 2348 (1998): "TFTP Blocksize Option".
- [41] IETF RFC 2460 (1998): "Internet Protocol, Version 6 (IPv6), Specification".
- [42] IETF RFC 2461 (1998): "Neighbor Discovery for IP Version 6 (IPv6)".

- [43] IETF RFC 2474 (1998): "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers".
- [44] IETF RFC 2616 (1999): "Hypertext Transfer Protocol -- HTTP/1.1".
- [45] IETF RFC 4639 (2006): "Cable Device Management Information Base for Data-Over-Cable Service Interface Specification (DOCSIS) Compliant Cable Modems and Cable Modem Termination Systems".
- [46] IETF RFC 2710 (1999): "Multicast Listener Discovery (MLD) for IPv6".
- [47] IETF RFC 2786 (2000): "Diffie-Helman USM Key Management Information Base and Textual Convention".
- [48] IETF RFC 3046 (2001): "DHCP Relay Agent Information Option".
- [49] IETF RFC 3203 (2001): "DHCP reconfigure extension".
- [50] IETF RFC 3256 (2002): "The DOCSIS (Data-Over-Cable Service Interface Specifications) Device Class DHCP (Dynamic Host Configuration Protocol) Relay Agent Information Sub-option".
- [51] IETF RFC 3315 (2003): "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)".
- [52] IETF RFC 3376 (2002): "Internet Group Management Protocol, Version 3".
- [53] IETF RFC 3513 (2003): "Internet Protocol Version 6 (IPv6) Addressing Architecture".
- [54] IETF RFC 3810 (2004): "Multicast Listener Discovery Version 2 (MLDv2) for IPv6".
- [55] IETF RFC 4361 (2006): "Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)".
- [56] IETF RFC 4601 (2006): "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)".
- [57] IETF RFC 4605 (2006): "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding (IGMP/MLD Proxying)".
- [58] IETF RFC 4607 (2006): "Source-Specific Multicast for IP".
- [59] IETF RFC 4861 (2007): "Neighbor Discovery for IP version 6 (IPv6)".
- [60] IETF RFC 4862 (2007): "IPv6 Stateless Address Autoconfiguration".
- [61] IETF RFC 4649 (2006): "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Remote-ID Option".
- [62] IETF RFC 5460 (2009): "DHCPv6 Bulk Leasequery".
- [63] FIPS PUB 180-1 (April 1995): "Secure Hash Standard".
- [64] IEEE 802.1P: "Traffic Class Expediting and Dynamic Multicast Filtering".
- [65] IEEE 802.1Q: "Virtual LANs".
- [66] IEEE 802.2: "Standard for Information technology -- Telecommunications and information exchange between systems -- Local and metropolitan area networks -- Specific requirements -- Part 2: Logical Link Control".

2.2 Informative references

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

- [i.1] Cable Television Laboratories, Inc. SP-CMTS-NSI-I01-960702 (July 1996): "Data Over Cable Interface Specifications - Cable Modem Termination System - Network Side Interface Specification".
- [i.2] ITU-T Recommendation J.83 (2007), Annex A: "Digital multi-program systems for television, sound and data services for cable distribution".
- [i.3] ITU-T Recommendation J.83 (2007), Annex B: "Digital multi-program systems for television, sound and data services for cable distribution".
- [i.4] ITU-T Recommendation J.83 (2007), Annex C: "Digital multi-program systems for television, sound and data services for cable distribution".
- [i.5] IETF RFC 3168 (September 2001): "The Addition of Explicit Congestion Notification (ECN) to IP", K. Ramakrishnan, S. Floyd, D. Black.
- [i.6] IETF RFC 3260 (April 2002): "New Terminology and Clarifications for Diffserv", D. Grossman.
- [i.7] IETF RFC 4291 (February 2006): "IP Version 6 Addressing Architecture", R. Hinden, S. Deering.
- [i.8] ITU-T Recommendation G.711: "Pulse code modulation (PCM) of voice frequencies".
- [i.9] IETF RFC 791: "Internet Protocol".
- [i.10] IEEE 802.3: "IEEE Standard for Local and Metropolitan Area Networks - Specific requirements Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications".
- [i.11] IEEE 802.1: "Working Group".

3 Definitions and abbreviations

3.1 Definitions

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

active codes: set of spreading codes which carry information in an S-CDMA upstream. The complementary set, the unused codes, are idle and are not transmitted. Reducing the number of active codes below the maximum value of 128 may provide advantages including more robust operation in the presence of colored noise

Address Resolution Protocol (ARP): protocol of the IETF for converting network addresses to 48-bit Ethernet addresses

Advanced Time Division Multiple Access (A-TDMA): DOCSIS 3.0 TDMA mode (as distinguished from DOCSIS 1.x TDMA)

allocation: group of contiguous mini-slots in a MAP which constitute a single transmit opportunity

American National Standards Institute (ANSI): U.S. standards body

bandwidth allocation map: MAC Management Message that the CMTS uses to allocate transmission opportunities to CMs

BITS encoding: octet string using a BITS encoding represents a zero-indexed linear array of 8*N bits, with the most significant bit of each byte representing the lowest-indexed bit. Bit positions increase from left to right. For example, bit position 0 is the most significant bit of the most significant (leftmost) byte, encoded as hex 0x80. Unspecified bit positions are assumed as zero. Unimplemented bit positions are ignored

bonded channel set: identified set of upstream or downstream channels among which a stream of packets is distributed

bonding group: list of channels providing a means to identify the specific channels bonded together

bridged network: set of- IEEE 802 LANs interconnected by IEEE 802.1D [16] MAC bridges

bridging CMTS: CMTS that makes traffic forwarding decisions between its Network System Interfaces and MAC Domain Interfaces based upon the Layer 2 Ethernet MAC address of a data frame

burst: single continuous RF signal from the upstream transmitter, from transmitter on to transmitter off

Byte: contiguous sequence of eight bits (an octet)

Cable Modem (CM): modulator-demodulator at subscriber locations intended for use in conveying data communications on a cable television system

cable modem service froup: HFC plant topology, the complete set of downstream and upstream channels within a single CMTS that a single Cable Modem could potentially receive or transmit on. In most HFC deployments, a CM-SG corresponds to a single Fibre Node. Usually, a CM-SG serves multiple CMs

Cable Modem Termination System (CMTS): Cable Modem termination system, located at the cable television system head-end 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 (NSI): network side interface, defined in [i.1], between a CMTS and the equipment on its network side

Cable Modem to CPE interface: interface, defined in [3], between a CM and CPE

capture bandwidth: sum of the Tuning Bands in the TB List in MHz

ceiling (ceil): mathematical function that returns the lowest-valued integer that is greater than or equal to a given value

channel: See Radio Frequency Channel.

channel bonding: logical process that combines the data packets received on multiple independent channels into one higher-speed data stream. Channel bonding can be implemented independently on upstream channels or downstream channels

chip: each of the 128 bits comprising the S-CDMA spreading codes

classifier: 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

CMCI port: physical interface of the CM to which a CPE device can attach

codeword: element of an error-correcting code used to detect and correct transmission errors

Continuous Concatenation and Fragmentation (CCF): method of packing data into segments for upstream transmission in Multiple Transmit Channel Mode

Converged Interconnect Network (CIN): network (generally gigabit Ethernet) that connects an M-CMTS Core to an EQAM

CPE interface: interface that is either a CMCI Port or a Logical CPE interface

Customer Premises Equipment (CPE): equipment at the end user's premises; may be provided by the end user or the service provider

Customer premises equipment Controlled Cable Modem (CCCM): CPE Controlled Cable Modem. Refer to [3].

Data Link Layer (DLL): layer 2 in the Open System Interconnection (OSI) architecture; the layer that provides services to transfer data over the transmission link between open systems

data rate: throughput, data transmitted in units of time usually in bits per second (bps)

decibels (dB): unit to measure the relative levels of current, voltage or power. An increase of 3 dB indicates a doubling of power, an increase of 10 dB indicates a 10x increase in power and an increase of 20 dB indicates a 100x increase in power

Decibel-Millivolt: dB measurement system wherein 0 dBmV is defined as 1 millivolt over 75 Ω

DOCSIS 1.x: abbreviation for "DOCSIS 1.0 or 1.1"

DOCSIS 2.0 Mode: CM operates in this mode when:

- 1) Multiple Transmit Channel (MTC) Mode is disabled;
- 2) the Enable 2.0 Mode configuration setting in the REG-RSP is set to 1 (Enable) explicitly or by default; and
- 3) it operates on an upstream channel using the burst descriptors associated with IUC 9, 10 and 11 as opposed to IUC 5 and 6.

A CM is enabled for DOCSIS 2.0 Mode when the Enable 2.0 Mode configuration setting in the REG-RSP is set to 1 (Enable). A CM may be enabled for DOCSIS 2.0 Mode but may not be operating in DOCSIS 2.0 Mode. When a CM has MTC Mode enabled, the CM is not considered to be in DOCSIS 2.0 Mode even if some of the upstream channels it is using are operating with post-1.1 DOCSIS physical layer mechanisms. Therefore, "DOCSIS 2.0 Mode" does not have relevance for a CM operating in MTC Mode.

downstream: in cable television, the direction of transmission from the head-end to the subscriber

downstream bonded service flow: downstream Service Flow assigned to a Downstream Bonding Group

Downstream Bonding Group (DBG): subcomponent object of a MAC Domain that distributes packets from an assigned set of Downstream Bonding Service Flows to an associated set of Downstream Channels of that MAC Domain

Downstream Channel (DC): physical layer characteristics and MAC layer parameters and functions associated to a DOCSIS forward channel

Downstream Channel Identifier (DCID): 8-bit identifier that distinguishes a Downstream Channel within a MAC Domain. DCID values may be assigned locally by the CMTS or externally by CMTS configuration

Downstream Interface (DI): as a term, refers to either a Downstream Channel (DC) or a Downstream Bonding Group (DBG). A DI is not a separate object in the object model

Downstream M-CMTS Channel: object representing the M-CMTS DEPI session (see [2]) that carries the DOCSIS MAC-Layer contents of a single Downstream RF Channel

Downstream RF Channel: CMTS object representing the physical transmission of the MAC-Layer contents of a DOCSIS downstream RF signal at a single center frequency. A DRF object implements the functions of:

- FEC Encoding;
- MPEG2 Convergence;
- QAM modulation; and
- Physical RF transmission.

downstream service extended header: DOCSIS extended header that contains a Downstream Service ID (DSID)

Downstream Service Group: complete set of Downstream Channels (DCs) from a single CMTS that could potentially reach a single Cable Modem. A DS-SG corresponds to a broadband forward carrier path signal from one CMTS. In an HFC deployment, a DS-SG corresponds to the downstream fibre transmission from one CMTS to one or more Fibre Nodes

Downstream Service Identifier (DSID): 20-bit value in a DOCSIS extended header that identifies a stream of packets distributed to the same cable modem or group of cable modems. The DSID value is unique within a MAC Domain. For sequenced packets, the DSID identifies the resequencing context for downstream packet bonding in the CM

DSID-indexed Payload Header Suppression: method used to provide Payload Header Suppression to downstream multicast sessions. The DSID identifies the PHS rule. The DSID, PHSS and PHSF are signaled as part of the multicast DSID encodings in a REG-RSP or DBC-REQ

Dual Stack Management: See Dual Stack Management Mode.

Dual Stack Management Mode: mode of DOCSIS cable modem operation in which the modem is manageable simultaneously via both IPv4 and IPv6 addresses

Duplicate Address Detection: Defined in the RFC 4605 [57].

Dynamic Host Configuration Protocol: Internet protocol used for assigning network-layer (IP) addresses

dynamic range: 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

Edge Quadrature Amplitude Modulator (EQAM): in the M-CMTS architecture, a network element that terminates DEPI sessions and implements the physical Downstream RF Channel for those sessions. The EQAM terminates Downstream M-CMTS Channels and forwards their DOCSIS MAC-Layer contents to Downstream RF Channels

egress interface: CPE interface through which the cable modem transmit traffic

end user: human being organization or telecommunications system that accesses the network in order to communicate via the services provided by the network

epoch time: time elapsed since 1 January 1970 00:00:00. This is usually expressed in seconds

fibre node: in HFC, a point of interface between a fibre trunk and the coaxial distribution

flooding: operation of an L2 Bridge in which it replicates an L2PDU addressed to a group MAC or unlearned individual MAC address to all Bridge Ports other than the L2PDU's ingress port

floor: mathematical function that returns the highest-valued integer that is less than or equal to a given value

Forward Error Correction (FEC): enables the receiver to detect and fix errors to packets without the need for the transmitter to retransmit packets

Frame: See MAC frame, S-CDMA frame and MPEG frame.

group delay: difference in transmission time between the highest and lowest of several frequencies through a device, circuit or system

Group Service Flow (GSF): downstream Service Flow for packets forwarded to hosts reached through a group of Cable Modems. A GSF may be either a Bonded GSF (B-GSF) or a Non-Bonded GSF (NB-GSF)

guard time: measured in modulation symbols, 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

head-end: central location on the cable network that is responsible for injecting broadcast video and other signals in the downstream direction

header: protocol control information located at the beginning of a protocol data unit

Hertz: unit of frequency equivalent to one cycle per second

NOTE: See also kiloHertz (kHz) and MegaHertz (MHz).

Hybrid Fibre/Coaxial System: broadband bidirectional shared-media transmission system using fibre trunks between the head-end and the fibre nodes and coaxial distribution from the fibre nodes to the customer locations

Individual MAC Address: IEEE 6-byte MAC address with the first transmitted bit (the group bit) set to 0, indicating that the address refers to a single MAC host. For the Ethernet MAC addresses of DOCSIS, the group bit is the least significant bit of the first byte of the MAC address

Individual Service Flow (ISF): Downstream Service Flow for packets forwarded to hosts reached through an individual Cable Modem. An ISF may be either a Bonded ISF(B-ISF) or a Non-Bonded ISF (NB-ISF)

information element: fields that make up a MAP and define individual grants, deferred grants, etc.

Institute of Electrical and Electronics Engineers (IEEE): voluntary organization which, among other things, sponsors standards committees and is accredited by the American National Standards Institute

integrated Cable Modem Termination System (CMTS): CMTS wherein all components are integrated into a single chassis as opposed to a modular CMTS

International Electrotechnical Commission (IEC): international standards body

Internet Control Message Protocol (ICMP): Internet network-layer protocol

Internet Engineering Task Force (IETF): body responsible, among other things, for developing standards used in the Internet

Internet Group Management Protocol (IGMP): network-layer protocol for managing multicast groups on the Internet

Internet Protocol (IP): computer network protocol (analogous to written and verbal languages) that all machines on the Internet must know so that they can communicate with one another. IP is a layer 3 (network layer) protocol in the OSI model. The vast majority of IP devices today support IP version 4 (IPv4) defined in RFC 791 [i.9], although support for IP version 6 (IPv6, RFC 2460 [41]) is increasing

Interval Usage Code (IUC): field in MAPs and UCDs to link burst profiles to grants

latency: time taken for a signal element to pass through a device

layer: subdivision of the Open System Interconnection (OSI) architecture, constituted by subsystems of the same rank

Layer 2 Protocol Data Unit: sequence of bytes consisting of a destination MAC address, a source MAC address, optional Tag Headers, Ethertype/Length, L2 Payload and CRC

Layer 2 Virtual Private Network: set of LANs and the L2 Forwarders between them that enable hosts attached to the LANs to communicate with L2PDUs. A single L2VPN forwarding L2PDUs based only on the Destination MAC address of the L2PDU, transparent to any IP or other Layer 3 address. A cable operator administration domain supports multiple L2VPNs, one for each subscriber enterprise to which Transparent LAN Service is offered

learning: operation of a layer 2 Bridge by which it associates the Source MAC address of an incoming L2PDU with the bridge port from which it arrived

link layer: See Data Link Layer.

Load Balancing Group (LBG): full or partial subset of a MAC Domain Cable Modem Service Group (MD-CM-SG) to which a CM is administratively assigned. LBGs contain at least one upstream channel and at least one downstream channel

Local Area Network (LAN): non-public data network in which serial transmission is used for direct data communication among data stations located on the user's premises

local log: volatile or nonvolatile log stored within a network element

Logical CPE Interface: logical interface between the embedded cable modem and an eSAFE

Logical (Upstream) Channel: MAC entity identified by a unique channel ID and for which bandwidth is allocated by an associated MAP message. A physical upstream channel may support multiple logical upstream channels. The associated UCD and MAP messages completely describe the logical channel

Logical Link Control (LLC): sub-layer of the second layer (Data Link Layer) in the Open Systems Interconnection seven-layer model for communications protocols standardized by the International Organization for Standardization (ISO), that is responsible for multiplexing transmitted messages, demultiplexing received messages and providing message flow control

MAC Domain: subcomponent of the CMTS that provides data forwarding services to a set of downstream and upstream channels

MAC Domain Cable Modem Service Group: subset of a CM-SG which is confined to the DCs and UCs of a single MAC domain. An MD-CM-SG differs from a CM-SG only if multiple MAC domains are assigned to the same CM-SGs

MAC Domain Downstream Service Group: subset of a Downstream Service Group (DS-SG) which is confined to the Downstream Channels of a single MAC domain. An MD-DS-SG differs from a DS-SG only when multiple MAC domains are configured per DS-SG

MAC Domain Interface: interface of a MAC Domain to a CMTS Forwarder

MAC Domain Upstream Service Group: subset of an Upstream Service Group (US-SG) which is confined to the Upstream Channels of a single MAC Domain. An MD-US-SG differs from a US-SG only when multiple MAC domains are defined per US-SG

MAC Frame: MAC header plus optional protocol data unit

MAP: See Bandwidth Allocation Map.

MDF-capable CM: CM that reports an MDF capability of 1 or 2 in the Modem Capabilities encoding

MDF-disabled: MDF-capable CM is said to be MDF-disabled when the CMTS sets the value of 0 for the MDF capability in the Modem Capabilities encoding of the REG-RSP(-MP)

MDF-enabled: CM is said to be MDF-enabled when the CMTS returns the value of 1 or 2 for the MDF capability in the Modem Capabilities encoding of the REG-RSP(-MP)

MDF-incapable CM: CM that reports an MDF capability of 0 or does not report an MDF capability in the Modem Capabilities encoding

Media Access Control: 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

Media Access Control Address: hardware address of a device connected to a shared medium

Media Access Control Frame: MAC header plus optional PDU

Media Access Control sublayer: sub-layer of the second layer (Data Link Layer) in the Open Systems Interconnection sublayer seven-layer model for communications protocols standardized by the International Organization for Standardization (ISO), that is responsible for determining which transmitter is allowed access to the communication medium and uses the services of the Physical Layer to provide services to the Logical Link Control (LLC) sublayer

MegaHertz: one million cycles per second

Micro-reflections: echoes in the forward transmission path due to impedance mismatches between the physical plant components. Micro-reflections are distinguished from discrete echoes by having a time difference (between the main signal and the echo) on the order of 1 microsecond. Micro-reflections cause departures from ideal amplitude and phase characteristics for the transmission channel

Microsecond (μ s): one millionth of a second

Millisecond (ms): one thousandth of a second

mini-slot: "mini-slot" is an integer multiple of 6,25-microsecond increments

Modular Cable Modem Termination System: CMTS composed of discrete functional blocks linked together using Gigabit Ethernet links

modulation rate: signalling rate of the upstream modulator (1 280 kHz to 5 120 kHz). In S-CDMA, the chip rate. In TDMA, the channel symbol rate

Moving Picture Experts Group (MPEG): voluntary body which develops standards for digital compressed moving pictures and associated audio

multicast client: entity with a unique MAC address that receives multicast packets

Multicast Downstream Service Identifier Forwarding Capable Cable Modem: cable modem that reports a nonzero value for the Multicast DSID Forwarding capability in the REG-REQ message

Multiple Outstanding Requests: ability of the cable modem to make additional bandwidth request for new packets for a service flow while one or more previous requests for older packets remain unfulfilled

Multiple System Operator: corporate entity that owns and/or operates more than one cable system

Multiple Transmit Channel Mode: upstream operation of the cable modem and cable modem termination system MAC layer using continuous concatenation and fragmentation to segment traffic and queue-depth based requesting

nanosecond (ns): one billionth of a second

network layer: layer 3 in the Open Systems Interconnection (OSI) architecture; the layer that provides services to establish a path between open systems

network management: functions related to the management of data link layer and physical layer resources and their stations across the data network supported by the hybrid fibre/coaxial system

Non-bonded Service Flow: Service Flow assigned to a single channel, rather than a Bonding Group

Non-primary Downstream Channel: Downstream Channel received by a cable modem which is not its Primary Downstream Channel

notification: information emitted by a managed object relating to an event that has occurred within the managed object

Open Systems Interconnection (OSI): 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

packet identifier: unique integer value used to identify elementary streams of a program in a single or multi -program MPEG-2 stream

partial service: modem is in a partial service mode of operation any time it is operating with a subset of the channels in the RCS and/or TCS because a channel has become unusable, either due to an inability to acquire a channel or because communication on a channel was lost during normal operation

Payload Header Suppression: transmitting or forwarding the data payload of a DOCSIS MAC frame without including header fields of the various protocol layers above the DOCSIS MAC layer. Suppression of header fields is selectable in DOCSIS

Physical 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 Sublayer: 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

Pre-3.0 DOCSIS: versions of Data-Over-Cable-Service-Interface-Specifications (DOCSIS) specifications prior to the DOCSIS 3.0 suite of specifications

Primary-Capable Downstream Channel: Downstream Channel which carries SYNC messages, MDD messages containing ambiguity resolution TLVs, as well as UCD and MAP messages for at least one upstream channel in each of the MD-CM-SG that the downstream channel reaches

Primary Channel: See Primary Downstream Channel.

Primary Downstream Channel: Primary-Capable Downstream Channel on which the DOCSIS 3.0 CM has achieved SYNC lock and successfully received an MDD message containing ambiguity resolution TLVs

Primary Service Flow: first service flow, in each direction, defined in the CM configuration file

protocol: set of rules and formats that determines the communication behavior of layer entities in the performance of the layer functions

Quadrature Amplitude Modulation (QAM): method of modulating digital signals onto a radio-frequency carrier signal involving both amplitude and phase coding

QAM channel: analog RF channel that uses quadrature amplitude modulation (QAM) to convey information

Quadrature Phase Shift Keying: method of modulating digital signals onto a radio-frequency carrier signal using four phase states to code two digital bits

Quality of Service Parameter Set: set of Service Flow Encodings that describe the Quality of Service attributes of a Service Flow or a Service Class

Queue-depth Based Request: request in multiples of bytes based on the CM's queue depth and QoS parameters for a specific service flow. This request does not include any estimation of physical layer overhead

Radio Frequency (RF): in cable television systems, this refers to electromagnetic signals in the range 5 MHz to 1 000 MHz

radio frequency channel : frequency spectrum occupied by a signal. Usually specified by center frequency and bandwidth parameters

Radio Frequency Interface (RFI): term encompassing the downstream and the upstream radio frequency interfaces

Ranging SID: SID used for ranging on a specific channel

Receive Channel Configuration (RCC): CMTS send the RCC encoding in the REG-RSP message. The RCC contains TLVs to initially configure a CM's Receive Channels (RCs) and Receive Modules (RMs)

Receive Channel Profile (RCP): describes a logical representation of the CM's downstream physical layer in terms of Receive Channels (RCs) and Receive Modules (RMs). A Cable Modem reports its ability to receive multiple channels with one or more RCP Encodings in a REG-REQ-MP message

Receive Channel Set (RCS): set of downstream channels assigned to an individual CM is called its Receive Channel Set and is explicitly configured by the CMTS using the RCC encodings

Receive Module (RM): component in the CM physical layer implementation shared by multiple Receive Channels

Request for Comments (RFC): technical policy document of the IETF. These documents can be accessed on the World Wide Web at <http://www.rfc-editor.org/>

Resequencing Channel List: list of channels on which the CM receives packets labeled with that DSID

Resequencing Context: CM Resequencing Context, identified by a Resequencing DSID, is the set of Downstream Resequencing Channel List, Sequence Change Count and DSID Resequencing Wait Time. Downstream packets containing a Resequencing DSID and a sequence number are delivered, resequenced and forwarded according to the attributes of the Resequencing Context

Resequencing Downstream Service Identifier: downstream service identifier for which the CMTS signals packet resequencing attributes

Routing CMTS: CMTS that makes traffic forwarding decisions between its Network System Interfaces and MAC Domain Interfaces based upon the Layer 3 (network) address of a packet

S-CDMA Frame: two dimensional representation of mini-slots, where the dimensions are codes and time. An S-CDMA frame is composed of p active codes in the code dimension and K spreading intervals in the time dimension. Within the S-CDMA frame, the number of mini-slots is determined by the number of codes per mini-slot (c) and p , the number of active codes in the S-CDMA frame. Each S-CDMA frame thus contains s mini-slots, where $s = p/c$ and each mini-slot contains $c \cdot K$ information (QAM) symbols

Security Association: set of security information shared by two devices in order to support secure communications between the devices across a network

Security Association Identifier: 14-bit handle used to identify a Security Association between a CM and a CMTS

segment: contiguous burst of upstream data traffic (IUCs 5, 6, 9, 10 or 11) allocated using a single grant element in a MAP message

Segment Header ON: mode of Upstream DOCSISv3.0 Operation where segment headers are used for each segment. This mode is provisioned per upstream service flow

Segment Header OFF: mode of Upstream DOCSIS 3.0 Operation where segment headers are not used for any segment. This mode is provisioned per upstream service flow and prohibits fragmenting a packet across segment boundaries

Selectable Active Codes: methodology to determine the set of active codes and its complement, the set of unused codes. In SAC mode 1, a consecutive set of codes starting with code 0 are unused. In SAC mode 2, the active codes are selectable via a 128-bit string

Service Class: 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: ASCII string by which a Service Class may be referenced in modem configuration files and protocol exchanges

Service Flow (SF): MAC layer transport service which provides unidirectional transport of packets from the upper layer service entity to the RF and shapes, polices and prioritizes traffic according to QoS traffic parameters defined for the Flow

Service Flow Identifier (SFID): identifier assigned to a service flow by the CMTS [32 bits]

Service Group (SG): is formally defined as the complete set of upstream and downstream channels that can provide service to a single subscriber device. This includes channels from different DOCSIS MAC Domains and even different CMTSs as well as video EQAMs

Service Identifier (SID): 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 Cluster: group of SIDs containing one and only one SID for each upstream channel within an upstream bonding group and treated the same from a request/grant perspective

SID Cluster Group: set of all SID Clusters associated with a specific service flow

Simple Network Management Protocol (SNMP): network management protocol of the IETF

Spreader-Off S-CDMA Burst: transmission from a single CM in a spreader-off frame on an S-CDMA channel defined by the time in which the cable modem's transmitter turns on to the time it turns off. There will generally be several spreader off bursts in a spreader-off frame

Spreading Codes: set of 128 binary sequences of 128 bits each which may be used to carry information in the S-CDMA upstream. The spreading codes are orthogonal, meaning their cross-correlation is zero. Each code carries a single QAM symbol of information when the code's amplitude and phase are modulated

Spreading Interval: Time to transmit a single complete S-CDMA spreading code, equal to the time to transmit 128 chips. Also, time to transmit a single information (QAM) symbol on an S-CDMA channel

sublayer: subdivision of a layer in the Open System Interconnection (OSI) reference model

subnetwork: subnetworks are physically formed by connecting adjacent nodes with transmission links

subscriber: See end user.

subsystem: 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

Synchronous Code Division Multiple Access (SCDMA): multiple access physical layer technology in which different transmitters can share a channel simultaneously. The individual transmissions are kept distinct by assigning each transmission an orthogonal "code". Orthogonality is maintained by all transmitters being precisely synchronized with one another

syslog: protocol that provides the transport of event notification messages across IP networks

tag header: 16-bit Tag Protocol ID (0x8100) followed by a 16-bit Tag Control field. The Tag Control field consists of a 3-bit User Priority field, a 1-bit Canonical Format Indicator and a 12-bit VLAN ID [16]

tick: 6,25-microsecond time intervals that are the reference for upstream mini-slot definition and upstream transmission times

Time Division Multiple Access (TDMA): digital technology that enables a large number of users to access, in sequence, a single radio frequency channel without interference by allocating unique time slots to each user within each channel

traffic segmentation: dividing upstream traffic into one or more segments on one or more upstream channels

Transmission Control Protocol (TCP): transport-layer Internet protocol which ensures successful end-to-end delivery of data packets without error

Transmit Channel Configuration (TCC): TLV settings in Registration and DBC MAC Management Messages that define operations such as addition, deletion, change, replacement or re-ranging of one or more upstream channels in the Transmit Channel Set of a cable modem

Transmit Channel Set (TCS): set of upstream channels that a cable modem is configured to use for upstream transmission. Each upstream service flow of the cable modem may be associated with some or all of the channels in the Transmit Channel Set (TCS). The TCS of a cable modem is conveyed from a CMTS to a cable modem through the Transmit Channel Configuration (TCC) field in the Registration Response message

trap: unconfirmed SNMP message for asynchronous notification of events from an SNMP entity

Trivial File Transfer Protocol (TFTP): Internet protocol for transferring files without the requirement for user names and passwords that is typically used for automatic downloads of data and software

Type/Length/Value (TLV): 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 of the element

upstream: direction from the subscriber location toward the head-end

Upstream Bonding Group (UBG): subcomponent object of a MAC Domain that collects and resequences/reassembles Upstream Segments from a UBSF from an administered set of UCs

Upstream Bonded Service Flow: upstream Service Flow assigned to an Upstream Bonding Group

Upstream Channel: physical layer characteristics and MAC layer parameters and functions associated to a DOCSIS reverse channel

Upstream Channel Bonding: ability of the cable modem and cable modem termination system to support allocating traffic for a single Service Flow across two or more upstream channels

Upstream Channel Descriptor (UCD): MAC Management Message used to communicate the characteristics of the upstream physical layer to the cable modems

Upstream Channel Identifier (UCID): 8-bit identifier that distinguishes an Upstream Channel within a MAC Domain

Upstream Drop Classifier (UDC): set of matching criteria that the CM applies to each packet in order to determine whether to filter (drop) upstream traffic

Upstream Interface: term that refers to either an Upstream Channel or Upstream Bonding Group

Upstream Physical Channel: set of Upstream Channels received at the same Upstream RF Interface Port with overlapping frequency. Assigned ifType docsCableUpstream (129)

Upstream RF Interface Port: physical RF connector that receives multiple Upstream Physical Channels at different upstream frequencies

Upstream Service Group (US-SG): complete set of Upstream Channels (UCs) within a single CMTS potentially reachable by the transmission of a single Cable Modem. In an HFC deployment, a US-SG corresponds to the physical combining of the upstream reverse carrier path signal from one or more Fibre Nodes reaching a single CMTS

Virtual Local Area Network (VLAN): subset of the LANs of an IEEE 802.1 Bridged Network to which a VLAN Identifier (VLAN ID) is assigned. An L2VPN may consist of several VLANs, each with different VLAN IDs and even of VLANs on different IEEE 802.1 Bridged Networks with the same VLAN ID

3.2 Abbreviations

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

AD	Activity Detection
ANSI	American National Standards Institute
APM	Alternate Provisioning Mode
ARP	Address Resolution Protocol
ASCII	American Standard Code for Information Interchange
ASM	Any Source Multicast
ASN.1	Abstract Syntax Notation 1
A-TDMA	Advanced Time Division Multiple Access
ATM	Asynchronous Transfer Mode
BPI	Baseline Privacy Interface
BPI+	Baseline Privacy Interface Plus
BPKM	Baseline Privacy Key Management
CableLabs	Cable Television Laboratories, Inc.
CBR	Constant Bit Rate
CC	Confirmation Code
CCCM	Customer premises equipment Controlled Cable Modem
CCF	Continuous Concatenation and Fragmentation
CCITT	International Telegraph and Telephone Consultative Committee (see also ITU-T)
CIN	Converged Interconnect Network
CIR	Committed Information Rate
CM	Cable Modem
CMCI	Cable Modem to Customer Premises Equipment Interface
CMIM	Cable Modem Interface Mask
CM-SG	Cable Modem Service Group
CMTS	Cable Modem Termination System
CPE	Customer Premises Equipment
CRC	Cyclic Redundancy Check
CVC	Code Verification Certificate
CW	Continuous Wave
DA	Destination Address
DAD	Duplicate Address Detection
dB	Decibel
DBC	Dynamic Bonding Change
DBG	Downstream Bonding Group
DC	Downstream Channel
DCC	Dynamic Channel Change
DCD	Downstream Channel Descriptor
DCI	Device Class Identifier
DCID	Downstream Channel Identifier
DCS	Downstream Channel Set
DEPI	Downstream External-PHY Interface
DER	Distinguished Encoding Rules
DES	Data Encryption Standard
DHCP	Dynamic Host Configuration Protocol
DHCPv4	IPv4 version of the Dynamic Host Configuration Protocol
DHCPv6	IPv6 version of the Dynamic Host Configuration Protocol

DI	Downstream Interface
DIX	Digital Intel Xerox
DLL	Data Link Layer
DMAC	Destination Media Access Control address
DMPI	DOCSIS MAC-PHY Interface
DOCSIS®	Data-Over-Cable Service Interface Specifications
DPM	Dual-stack Provisioning Mode
DPV	DOCSIS Path Verify
DRFI	Downstream Radio Frequency Interface
DS	Downstream Services
DSA	Dynamic Service Addition
DSC	Dynamic Service Change
DSCP	Differentiated Services Code Point
DSD	Dynamic Service Deletion
DSG	DOCSIS Set-top Gateway
DSID	Downstream Service Identifier
DS-SG	Downstream Service Group
DTI	DOCSIS Timing Interface
DUID	DHCP Unique Identifier
DUT	Downstream Unencrypted Traffic
EAE	Early Authentication and Encryption
eCM	Embedded Cable Modem
EH	Extended Header
EHDR	Extended MAC Header
ELEN	Extended Header Length
EMA	Exponential Moving Average
eMTA	Embedded Media Transport Agent
ePS	Embedded Portal Services
EQAM	Edge QAM
eRouter	Embedded Router
eSAFE	Embedded Service/Application Functional Entity
EUI-64	64-bit Extended Unique Identifier
FC	Frame Control
FCRC	Fragment Cyclic Redundancy Check
FEC	Forward Error Correction
FHCS	Fragment Header Checksum
FIPS	Federal Information Processing Standard
FN	Fibre Node
FTP	File Transfer Protocol
GARP	Generic Attribute Registration Protocol
GCR	Group Classifier Rule
GLBG	General Load Balancing Group
GMAC	Group Media Access Control
GQC	Group QoS Configuration
GSF	Group Service Flow
HCS	Header Check Sequence
HFC	Hybrid Fibre/Coaxial
HMAC	Keyed-Hash Message Authentication Code
HTTP	HyperText Transfer Protocol
IC	Integrated Circuit
ICMP	Internet Control Message Protocol
I-CMTS	Integrated Cable Modem Termination System
IE	Information Element
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IP	Internet Protocol
IPDR	Internet Protocol Detail Record
IPTV	Internet Protocol Television
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6

IRT	Initial Retransmission Time
ISF	Individual Service Flow
ISO	International Standards Organization
ITU	International Telecommunications Union
ITU-T	Telecommunication Standardization Sector of the International Telecommunication Union
IUC	Interval Usage Code
kbps	Kilobits per second
L2	Layer 2
L2PDU	Layer 2 Protocol Data Unit
L2VPN	Layer 2 Virtual Private Network
LAN	Local Area Network
LBG	Load Balancing Group
LLC	Logical Link Control
LSB	Least Significant Bit
M/N	Relationship of integer numbers M,N that represents the ratio of the downstream symbol clock rate to the DOCSIS master clock rate
MAC	Media Access Control
Mbps	Megabits per second
M-CMTS	Modular Cable Modem Termination System
MD	Media Access Control Domain
MD-CM-SG	Media Access Control Domain Cable Modem Service Group
MDD	MAC Domain Descriptor
MD-DS-SG	Media Access Control Domain Downstream Service Group
MD-DS-SG-ID	Media Access Control Domain Downstream Service Group Identifier
MDF	Multicast DSID Forwarding
MD-US-SG	Media Access Control Domain Upstream Service Group
MD-US-SG-ID	Media Access Control Domain Upstream Service Group Identifier
MIB	Management Information Base
MIC	Message Integrity Check
MLD	Multicast Listener Discovery
MMM	MAC Management Message
MPEG	Moving Picture Experts Group
MRC	Maximum Retransmission Count
MRD	Maximum Retransmission Duration
MRT	Maximum Retransmission Time
MSAP	Media Access Control Service Access Point
MSB	Most Significant Bit
MSC	Maximum Scheduled Codes
MTA	Multimedia Terminal Adapter
MTC	Multiple Transmit Channel
MULPI	MAC and Upper Layer Protocols Interface
NACO	Network Access Control Object
ND	Neighbor Discovery
NSI	Network Side Interface
OID	Object Identifier
OP	Opportunity
OSI	Open Systems Interconnection
OSSI	Operations System Support Interface
OUI	Organizationally Unique Identifier
PDU	Protocol Data Unit
PER	Packet Error Rate
PHS	Payload Header Suppression
PHSF	Payload Header Suppression Field
PHSI	Payload Header Suppression Index
PHSM	Payload Header Suppression Mask
PHSR	Payload Header Suppression Rule
PHSS	Payload Header Suppression Size
PHSV	Payload Header Suppression Verify
PHY	Physical Layer
PID	Packet Identifier
PIM	Protocol Independent Multicast
PMD	Physical Media Dependent sublayer

POH	Physical Overhead
ppm	parts per million
PSP	Packet Streaming Protocol
PUSI	Payload Unit Start Indicator
QAM	Quadrature Amplitude Modulation
QI	Queue Indicator
QoS	Quality of Service
QPSK	Quadrature Phase Shift Keying
RA	Router Advertisement
RCC	Receive Channel Configuration
RCID	Receive Channel Identifier
RCP	Receive Channel Profile
RCP-ID	Receive Channel Profile Identifier
RCS	Receive Channel Set
RF	Radio Frequency
RFC	Request For Comments
RFI	Radio Frequency Interface
RM	Receive Module
RS	Router Solicitation
RSA	Rivest, Shamir, Adleman
RSVP	Resource Reservation Protocol
RTP	Real-time Transport Protocol
SA	Security Association
SA	Source Address
SAC	Selectable Active Codes
SAID	Security Association Identifier
SAV	Source Address Verification
SC	SID_Cluster
S-CDMA	Synchronous Code Division Multiple Access
SDL	Specification and Description Language
SF	Service Flow
SFID	Service Flow Identifier
SG	Service Group
SHA	Secure Hash Algorithm
SID	Service Identifier
SLAAC	Stateless Address Autoconfiguration
SM	Station Maintenance
SMA	Simple Moving Average
SNAP	SubNetwork Access Protocol
SNMP	Simple Network Management Protocol
SPI	Serial Peripheral Interface
SSM	Source Specific Multicast
STB	Set-Top Box
TCC	Transmit Channel Configuration
TCN	Topology Change Notification
TCP	Transmission Control Protocol
TCS	Transmit Channel Set
TDMA	Time Division Multiple Access
TEK	Traffic Encryption Key
TFTP	Trivial File Transfer Protocol
TLV	Type/Length/Value
ToD	Time of Day
TOS	Type of Service
TX	Transmit
UBG	Upstream Bonding Group
UCD	Upstream Channel Descriptor
UCID	Upstream Channel Identifier
UDC	Upstream Drop Classifier
UDP	User Datagram Protocol
UGS	Unsolicited Grant Service
UGSH	Unsolicited Grant Synchronization Header
URFI	Upstream RF Interface

US	Upstream
US-SG	Upstream Service Group
UTC	Coordinated Universal Time
VAD	Voice Activity Detection
VLAN	Virtual Local Area Network
VoIP	Voice over IP

4 Void

5 Overview and Theory of Operations

5.1 DOCSIS 3.0 MULPI Key Features

DOCSIS 3.0 introduces a number of features that build upon what was present in previous versions of DOCSIS. This specification includes the following key new features for the MAC and Upper Layer Protocols Interface.

- **Downstream Channel Bonding with Multiple Receive Channels:** DOCSIS 3.0 introduces the concept of a CM that receives simultaneously on multiple receive channels. Downstream Channel Bonding refers to the ability (at the MAC layer) to schedule packets for a single service flow across those multiple channels. Downstream Channel Bonding offers significant increases in the peak downstream data rate that can be provided to a single CM.
- **Upstream Channel Bonding with Multiple Transmit Channels:** DOCSIS 3.0 introduces the concept of a CM that transmits simultaneously on multiple transmit channels. Upstream Channel Bonding, refers to the ability to schedule the traffic for a single upstream service flow across those multiple channels. Upstream Channel Bonding offers significant increases in the peak upstream data rate that can be provided to a single CM. DOCSIS 3.0 also introduces other enhancements in the upstream request-grant process that improve the efficiency of the upstream link.
- **IPv6:** DOCSIS 3.0 introduces built-in support for the Internet Protocol version 6. DOCSIS 3.0 CMs can be provisioned with an IPv4 management address, an IPv6 management address or both. Further, DOCSIS 3.0 CMs can provide transparent IPv6 connectivity to devices behind the cable modem (CPEs), with full support for Quality of Service and filtering.
- **Source-Specific Multicast:** DOCSIS 3.0 supports delivery of Source-Specific IP Multicast streams to CPEs. Rather than extend the IP multicast protocol awareness of cable modems to support enhanced multicast control protocols, DOCSIS 3.0 takes a different approach. All awareness of IP multicast is moved to the CMTS and a new DOCSIS-specific layer 2 multicast control protocol between the CM and CMTS is defined which works in harmony with downstream channel bonding and allows efficient and extensible support for future multicast applications.
- **Multicast QoS:** DOCSIS 3.0 defines a standard mechanism for configuring the Quality of Service for IP multicast sessions. It introduces the concept of a "Group Service Flow" for multicast traffic that references a Service Class Name that defines the QoS parameters for the service flow.

5.2 Technical Overview

This specification defines the MAC layer protocols of DOCSIS 3.0 as well as requirements for upper layer protocols (e.g. IP, DHCP, etc.). DOCSIS 3.0 introduces new MAC layer features beyond what were present in earlier versions of DOCSIS.

Specifically, DOCSIS 3.0 defines a mechanism to increase the peak rate of upstream and downstream forwarding between the CMTS and a CM by utilizing multiple independent physical layer channels. This feature is termed channel bonding. Due to the inherent differences in the MAC layer definition for upstream transmission relative to downstream, the bonding mechanisms are themselves quite different in the two directions. This specification defines the requirements for CMs and CMTSs to support both upstream and downstream channel bonding.

DOCSIS 3.0 introduces a number of enhancements to the operation of upstream request and grant scheduling, including the ability to request in terms of bytes instead of mini-slots and to have multiple outstanding requests per upstream service flow. The set of upstream enhancements introduced with DOCSIS 3.0 is collectively called the "Multiple Transmit Channel Mode" of operation on the CM.

Additionally, DOCSIS 3.0 introduces enhancements to the way that IP multicast is handled. DOCSIS 1.1 and 2.0 required that cable modems actively participate in tracking layer-3 IP multicast group membership. DOCSIS 3.0, in contrast, provides a CMTS controlled layer-2 multicast forwarding mechanism. DOCSIS 3.0 also introduces the ability for cable operators to configure Quality of Service guarantees for multicast traffic. These features can be used to reliably deliver source-specific as well as any-source multicast sessions to clients behind the cable modem.

Finally, DOCSIS 3.0 introduces full support for IPv6, including the provisioning and management of a cable modem with an IPv6 address and the ability to manage and transport IPv6 traffic.

This specification also includes MAC layer protocol definitions for support of additional DOCSIS 3.0 features defined in the other DOCSIS 3.0 specifications: [15], [12] and [10].

5.2.1 CMTS and CM Models

5.2.1.1 CMTS Model

A CMTS is considered to be a DOCSIS network element that forwards packets between one or more Network Side Interface (NSI) ports (defined in [i.1]) and DOCSIS RF Interface (RFI) ports (defined in [4] and [12]). DOCSIS defines two types of CMTS:

- an "Integrated" CMTS that directly implements the NSI and RFI ports in a single network element; and
- a "Modular" CMTS that implements the NSI and Upstream RF Interfaces in a "Modular CMTS Core" network element and Downstream RF interfaces on an Edge QAM (EQAM) element.

This clause gives an overview of the CMTS model.

5.2.1.1.1 Types of CMTS

5.2.1.1.1.1 Integrated CMTS

An Integrated CMTS implements a single OSSI entity (SNMP agent, IPDR exporter) for Cable Operator configuration and management of the Downstream RF Interfaces (DRFIs) and Upstream RF Interfaces (URFIs) of the CMTS. Requirements for the DRFI are found in [4], requirements for the URFI are found in [12].

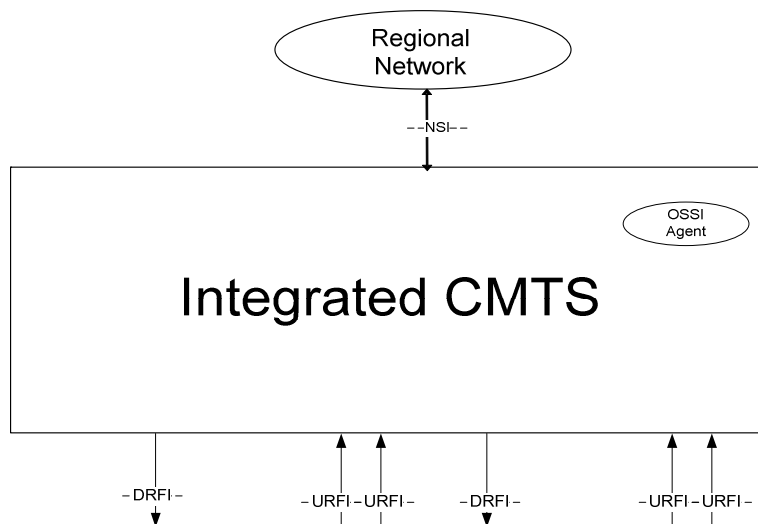


Figure 5-1: Integrated CMTS Network Diagram

An Integrated CMTS implements a single OSSI entity (SNMP agent, IPDR exporter) for Cable Operator configuration and management of the Downstream RF Interfaces [4] and Upstream RF Interfaces (URFIs) of the CMTS.

5.2.1.1.1.2 Modular CMTS

Figure 5-2 depicts a Modular CMTS (M-CMTS) network diagram. The M-CMTS Core implements the Network Side Interfaces and the Upstream RF Interfaces of a CMTS. The M-CMTS Core tunnels the contents of downstream DOCSIS channels across a Converged Interconnect Network (CIN) to one or more Edge QAMs (EQAMs) using the DOCSIS-standardized Downstream External Physical Interface [2]. The M-CMTS Core and all EQAMs are synchronized by a DOCSIS Timing Server using a standardized DOCSIS Timing Interface [6].

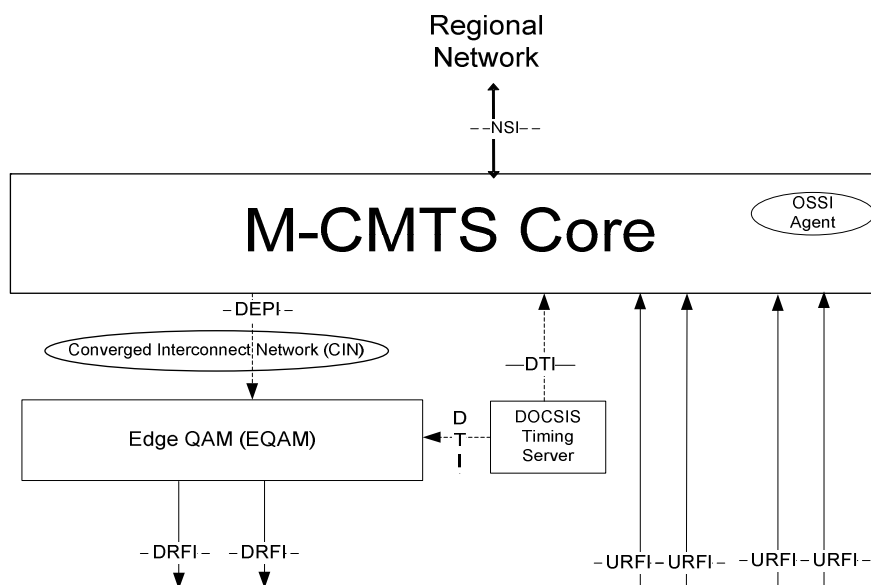


Figure 5-2: Modular CMTS Network Diagram

The only difference between the data forwarding models for an I-CMTS and an M-CMTS Core is how the contents of a downstream channel are transmitted. On an M-CMTS, the contents of a downstream channel are encapsulated into a DEPI Tunnel for transmission over the CIN to an EQAM, which are then modulated and transmitted by the Downstream RF port. In contrast, an I-CMTS the contents of a downstream channel are directly modulated and transmitted by the Downstream RF port.

In this specification, the term "CMTS" will refer to operation of both an Integrated CMTS and a Modular CMTS Core.

5.2.1.1.2 CMTS Internal Forwarding Model

Figure 5-3 depicts the logical operational model of internal packet forwarding within a CMTS.

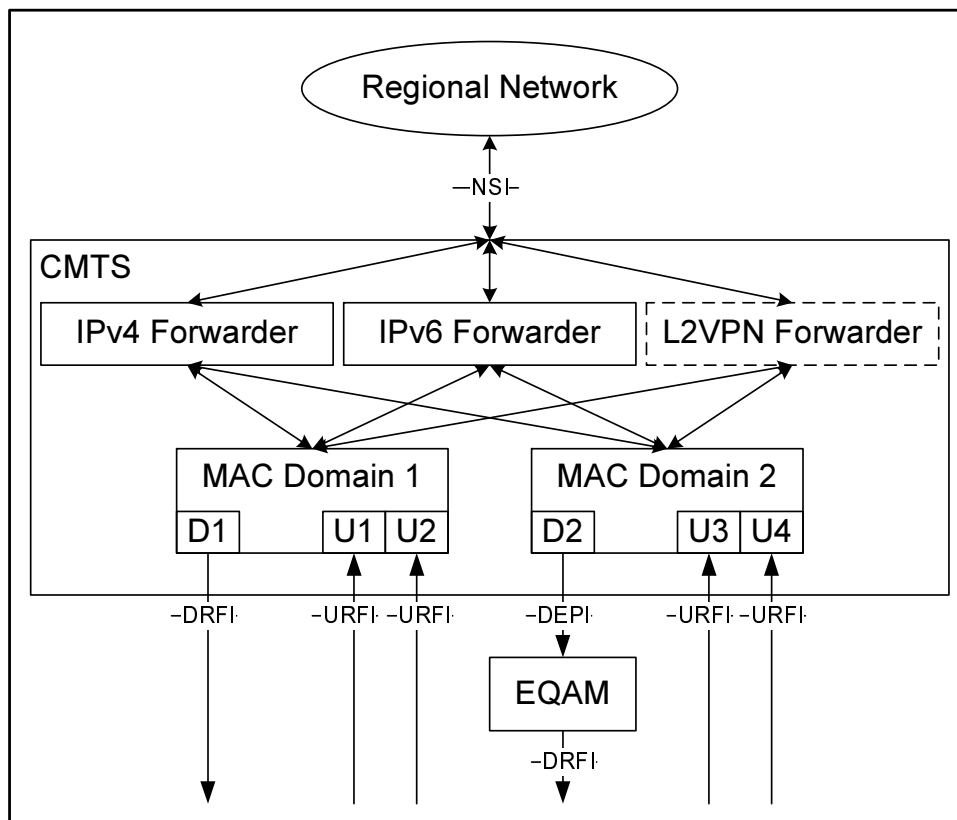


Figure 5-3: CMTS Internal Forwarding Model

The CMTS internal forwarding model consists of two types of sub-components:

- CMTS Forwarders which forward packets with layer 2 bridging or layer 3 routing; and
- MAC Domains which manage and forward data to and from cable modems reached by a set of downstream and upstream channels.

A CMTS Forwarder is responsible for forwarding packets between a Network Side Interface and the MAC Domains. In DOCSIS 3.0 the MAC Domain is not considered to forward data packets from its upstream to its own downstream channels; all upstream data packets are considered to be delivered to a CMTS Forwarder. DOCSIS 3.0 leaves most details of CMTS Forwarder operation to CMTS vendor-specific implementation. DOCSIS versions 1.0, 1.1 and 2.0 required that the CMTS permit IPv4 communication across the NSI port to CPE host(s) attached to CMs, along with IPv4 management of the CMTS and CMs themselves. DOCSIS 3.0 adds the requirement to manage CMs with IPv6, as well as to provide IPv6 connectivity across an NSI port to CPE IPv6 hosts. DOCSIS does not specify whether the CMTS implements layer 2 or layer 3 forwarding of the IPv4 and IPv6 protocols or prevent one protocol from being bridged and the other protocol from being routed. In addition, the DOCSIS Layer 2 Virtual Private Networking specification [8] standardizes transparent layer 2 forwarding between NSI ports and CM CPE interfaces and requires the implementation of an "L2VPN" CMTS Forwarder that is distinct from the "non-L2VPN" CMTS Forwarders for IPv4/IPv6 bridging or routing.

5.2.1.1.3 CMTS MAC Domain

A DOCSIS MAC Domain is a logical sub-component of a CMTS that is responsible for implementing all DOCSIS functions on a set of downstream channels and upstream channels. A CMTS MAC Domain contains at least one downstream channel and at least one upstream channel.

A MAC Domain is responsible for sending and receiving all MAC Management Messages (MMMs) to and from a set of CMs that are registered on that MAC Domain. A CM is registered to only a single MAC Domain at any given time.

A MAC Domain provides layer 2 data transmission services between the CMTS Forwarders and the set of CMs registered to that MAC Domain.

The MAC Domain classifies downstream packets into downstream "service flows" based on layer 2, 3 and 4 information in the packets. The MAC Domain schedules the packets for each downstream service flow to be transmitted on its set of downstream channels.

In the upstream direction, the MAC Domain indicates to a CMTS Forwarder component when a Layer 2 packet has been received from a particular CM. Each CMTS Forwarder component is responsible for forwarding and replicating (if necessary) Layer 2 packets between the MAC Domains and the NSI port(s) of a CMTS. All upstream DOCSIS Layer 2 packets are delivered to a CMTS Forwarder subcomponent; the MAC Domain does not directly forward Layer 2 packets from upstream to downstream channels. Since the CMTS Forwarder is responsible for building the Layer 2 Ethernet header of downstream Data PDU packets, the IPv4 ARP and IPv6 ND protocols are considered to be implemented within the CMTS Forwarder.

5.2.1.1.3.1 Downstream Data Forwarding in a MAC Domain

A MAC Domain provides downstream DOCSIS data forwarding service using the set of downstream channels associated with the MAC Domain. Each downstream channel in a MAC Domain is assigned an 8-bit Downstream Channel ID (DCID).

A downstream channel itself is defined as either:

- a "**Downstream (RF) Channel**", representing a single-channel downstream RF signal on a Downstream RF Port of an Integrated CMTS; or
- a "**Downstream M-CMTS Channel**", representing a single-channel downstream RF signal at a remote Edge QAM that is reached via a DEPI tunnel from an M-CMTS Core.

At an M-CMTS Core, the term "Downstream M-CMTS Channel" refers to the origination of a DEPI session. At an EQAM, the term "Downstream M-CMTS Channel" refers to the termination of a DEPI session.

5.2.1.1.3.2 Upstream Data Forwarding in a MAC Domain

An "upstream channel" can be used to refer to either:

- a "**Physical Upstream Channel**"; or
- a "**Logical Upstream Channel**" of a Physical Upstream Channel.

A "**Physical Upstream Channel**" is defined as the DOCSIS RF signal at a single center frequency in an upstream carrier path.

Multiple "**Logical Upstream Channels**" can share the center frequency of a Physical Upstream Channel, but operate in different subsets of the time domain. Transmit opportunities for each Logical Upstream Channel are independently scheduled by the CMTS.

A MAC Domain provides upstream DOCSIS data forwarding service using the set of logical upstream channels associated with the MAC Domain. Each logical upstream channel in a MAC Domain is assigned an 8-bit Upstream Channel ID (UCID).

All logical upstream channels operating at the same frequency on an Upstream RF Interface port are contained in the same MAC Domain.

5.2.1.2 CM Model

A CM is a DOCSIS network element that forwards (bridges) layer-2 traffic between a Radio Frequency Interface (RFI) and one or more Customer Premises Equipment ports.

5.2.2 DOCSIS MAC Operation

5.2.2.1 QoS

This clause provides an overview of the QoS protocol mechanisms and their part in providing end-to-end QoS.

Some of the Quality of Service related features described in this specification include:

- Packet Classification and Flow Identification.
- Service Flow QoS Scheduling with a set of QoS Parameters, including:
 - Traffic Priority.
 - Token Bucket Rate Shaping/Limiting.
 - Reserved (Guaranteed) Data Rate.
 - Latency and Jitter Guarantees.
- Both Static and Dynamic QoS Establishment.
- Two-Phase Activation Model for Dynamic QoS.

The majority of the QoS features in this specification were originally defined in [13]. This version of DOCSIS introduces a new feature to control prioritized data forwarding through the CM. This version of DOCSIS also defines a mechanism to configure QoS for downstream multicast traffic.

The various DOCSIS protocol mechanisms described in the present document can be used to support Quality of Service (QoS) for both upstream and downstream traffic through the CM and the CMTS.

The principal mechanism for providing QoS is to classify packets traversing the DOCSIS RF interface into a Service Flow and then to schedule those Service Flows according to a set of QoS parameters. A Service Flow is a unidirectional flow of packets that is provided a particular Quality of Service.

The requirements for Quality of Service include:

- A configuration and registration function for pre-configuring per-CM QoS Service Flows and traffic parameters.
- A signalling function for dynamically establishing QoS-enabled Service Flows and traffic parameters.
- CMTS MAC scheduling of downstream and upstream Service Flows based on QoS parameters for the Service Flow.
- CM and CMTS traffic-shaping, traffic-policing and traffic-prioritization based on QoS parameters for the Service Flow.
- 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.
- Assignment of Service Flows to particular upstream or downstream channels that reach the CM based on elements of the QoS parameter set 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. Specifically, the following behaviors are required in DOCSIS 3.0:

- In the upstream and downstream direction the CMTS can be configured to overwrite the DiffServ Field setting.
- The queuing of downstream PDU packets may be prioritized at the CMCI output of the CM by the Traffic Priority.

Additional behaviors are permitted, for example:

- The queuing of packets at the CMTS in the upstream and downstream directions may be based on the DiffServ Field.
- Downstream packets 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 are also assigned a 14-bit Service Identifier (SID) or one or more SID Clusters (which comprise a SID Cluster Group).

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. Similarly, the first downstream Service Flow describes the **Primary Downstream Service Flow**, which is the default Service Flow in the downstream direction. Additional Service Flows can be defined in the Configuration file to provide additional QoS services.

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.

5.2.2.1.1 Individual and Group Service Flows

Downstream Service Flows may be distinguished by whether they provide service to an individual CM or a group of CMs:

- Individual Service Flows are defined as Service Flows created by the Registration process of a single CM or a Dynamic Service Addition process to a single CM.
- Group Service Flows are created by the CMTS and may or may not be communicated to the CM.

A CMTS classifies packets offered for forwarding by an individual CM to an Individual Service Flow.

Individual Service Flows (and their classifiers) apply to only packets forwarded by the CMTS to hosts (embedded or non-embedded) reachable through a single CM. Individual Service Flow traffic is usually addressed to a unicast Destination MAC Address learned by the CMTS as reachable through that CM. Note, however, that with Layer 2 Virtual Private Network service [8], traffic with a non-unicast Destination MAC Address will also be forwarded through a single CM by requiring such traffic to be encrypted in the BPI Primary SAID of the CM.

Group Service Flows are intended primarily for traffic with a non-unicast Destination MAC Address, such as ARP broadcasts and downstream IP multicasts. A CMTS could send a downstream packet with a unicast Destination MAC Address on a Group Service Flow. One example is when the CMTS does not know to which CM the single Destination MAC Address is attached.

5.2.2.2 Channel Bonding

5.2.2.2.1 Downstream Channel Bonding

In order to provide peak downstream data rates in excess of 100 Mbps to customers, while maintaining interoperability with legacy CMs, DOCSIS 3.0 introduces a mechanism by which the CMTS dynamically distributes downstream packets over a *set* of downstream channels for delivery to a single CM. Each downstream channel in the set is a 6 MHz or 8 MHz (depending on region) MPEG Transport channel, consistent with those used in previous versions of DOCSIS. Each packet is tagged with a sequence number so that proper data sequencing is not lost if there are differences in latency between the channels in the set. The CM, in turn, has multiple receivers and is tuned to receive all of the channels in the set. The CM re-sequences the downstream data stream to restore the original packet sequence before forwarding the packets to its CPE port(s).

The term "downstream channel bonding" means the distribution of packets from the same service flow over different downstream channels. A "Downstream Bonding Group" (DBG) refers to the group of Downstream Channels over which the CMTS distributes the packets of a downstream service flow. The term "Downstream Bonding Group" is intended to refer to a set of two or more downstream channels, although during transition periods only a single channel may be defined or operational in a Downstream Bonding Group. Downstream Bonding Groups may either be statically provisioned by an operator or dynamically determined by the CMTS and need not be composed of adjacent RF channels.

In typical deployments there will be a large number of CMs tuned to the same Downstream Bonding Group. By distributing the downstream data traffic dynamically across the channels of that Bonding Group, the CMTS can ensure that the maximum gains from statistical multiplexing are achieved.

It is expected that deployments may have several downstream channels reaching a fibre node and that multiple (possibly overlapping) Downstream Bonding Groups will be defined, with CMs tuned to one or more of these Bonding Groups.

Further, each of the downstream channels in the set is capable of being configured to simultaneously support legacy CMs. The population of legacy CMs on a particular fibre node can then be dynamically balanced across the Downstream Bonding Group with each CM receiving a single channel at a time, in order to maintain the best service quality.

The CMTS is said to "assign" a downstream Service Flow to either a single downstream channel or to a Downstream Bonding Group. A Cable Operator can control the assignment of service flows to with a flexible "attribute" based assignment algorithm that is described in clause 8.1.1.

The term "Downstream Channel Set" (DCS) applies only in the CMTS and refers to an identified set of one or more channels over which packets of a service flow are scheduled. A DCS is either a single Downstream Channel or a multiple-channel Downstream Bonding Group. Each DCS to which the CMTS schedules packets is assigned a 16-bit Downstream Channel Set ID (DCS ID) by the CMTS. So a downstream Service Flow is considered to be "assigned" to a single DCS at any given point in time. A downstream Service Flow assigned to a DCS representing the multiple channels of a DBG is called a "bonded" downstream service flow. A downstream Service flow assigned to a DCS consisting of a single downstream channel is called a "non-bonded" Service Flow.

Because different downstream channels can have different latencies to the CM, packets of a bonded service flow distributed simultaneously across multiple channels can arrive at the CM out of order. DOCSIS 3.0 introduces the concept of a "packet sequence number" that is added to the frames of packets distributed over multiple channels. The packet sequence number is included in the 5-byte length version of a new Downstream Service Extended header (DS-EHDR) defined for DOCSIS 3.0. Downstream frames that include the 5-byte DS-EHDR are called "sequenced" frames.

A CM is expected to resequence only the frames that it will forward to CPEs; the CM does not resequence all packets transmitted downstream on a bonding group. Accordingly, a separate packet sequence number space is required for each individual CM that receives sequenced packets and indeed for each unique set of CMs receiving the sequenced frames of a multicast session.

A downstream sequence of packets is identified at the CMTS and CM by a 20-bit "Downstream Service ID" (DSID). The DSID identifies the CM or set of CMs intended to receive a downstream sequenced packet stream. The CMTS inserts a 5-byte Downstream Service Extended Header (DS EHDR) on each sequenced downstream packet to provide the DSID value and the packet's sequence number specific to that DSID. The use of a DSID to identify a particular packet stream sequence allows DOCSIS 3.0 CMs to filter downstream packets based on the DSID value and resequence only those packets intended to be forwarded through the CM.

The particular set of downstream channels on which a CM receives, distributed sequenced packets with a DSID label is called the Resequencing Channel Set of the DSID at that CM.

The stream of packets identified by a DSID is independent of a CMTS service flow. For example, the CMTS may utilize a single sequence number space (and one DSID) for one or more Service Flows forwarded to the same CM. Alternatively, the CMTS may classify different IP multicast sessions to the same Group Service Flow, in which case packets transmitted from the same group service flow could be transmitted with different DSIDs.

The set of downstream channels assigned to an individual CM is called its Receive Channel Set and is explicitly configured by the CMTS. The CMTS assigns a CM's bonded service flows to Downstream Bonding Groups that have channels in the CM's Receive Channel Set.

The CMTS assigns a Receive Channel Set to a CM by sending the CM a Receive Channel Configuration. The Receive Channel Set is the complete list of Downstream Channels that were defined in the Receive Channel Configuration.

The CMTS controls the Receive Channel Set for each CM and in doing so, can optimally support deployments where the aggregate data capacity needed (in terms of numbers of downstream channels) exceeds the number of channels that a single CM can receive. In this situation, the CMs can be dynamically balanced across the available downstream channels by manipulation of their respective Receive Channel Sets. For example, a particular fibre node could be configured to carry six downstream channels, yet each individual CM might only have the capability to receive four downstream channels simultaneously. By dynamically balancing the load (via Receive Channel Set assignments), the CMTS can provide the aggregate data capacity of all 6 downstream channels.

To support future CM hardware designs and limitations, DOCSIS 3.0 provides a flexible means for a CM to advertise its receiver characteristics (Receive Channel Profiles) and any limitations on Receive Channel Set assignment.

5.2.2.2.2 Upstream Channel Bonding

Cable operators would like to be able to provide higher upstream bandwidth per user in order to compete with FTTx offerings and provide services to small businesses.

The Cable Operators have stated an objective of 100 Mbps upstream throughput from a single user or group of users. Given the current impracticality of using very high orders of modulation (e.g. 1 024-QAM) and wider channels in the upstream, the only way to achieve the desired throughput using cable is to allow a user to transmit on multiple upstream channels simultaneously. This concept of a CM transmitting on multiple upstream channels simultaneously is new to DOCSIS and is referred to as Upstream Channel Bonding in that the smaller bandwidth upstream channels can be bonded together to create a larger bandwidth pipe.

The actual bonding process is controlled by the CMTS as part of the scheduling process via grants. The CM makes a request for bandwidth for a given service flow on one of the service flow's associated upstream channels. The CMTS then chooses whether to grant the request on one or more of the channels associated with that service flow. The CMTS is responsible for allocating the bandwidth across the individual upstream channels. This centralized control allows the system the best statistical multiplexing possible and allows the CMTS to do real-time load balancing of the upstream channels within a bonding group. When the CM receives grants over multiple channels, it divides its transmission according to the transmit time for each grant and the size of each grant. The CM places an incrementing sequence number in the traffic transmitted in each grant. The grants may be staggered in time across any or all of the channels and may require the CM to transmit on all bonded upstream channels simultaneously. The CMTS then uses the sequence number in the traffic to reconstruct the original data stream.

This mechanism for upstream channel bonding requires that the upstream channels be synchronized to a master clock source as discussed in clause 7.1. This synchronization requirement simplifies the clock domains and timing recovery in the CM. Other than this synchronization requirement, no other requirements are placed on the physical layer parameters of any of the channels within the Upstream Bonding Group. The individual channels can be any mix of modulation types, symbol rates, TDMA or S-CDMA as specified in the DOCSIS 3.0 Physical Layer specification [12] and can be any mix of adjacent or non-adjacent upstream channels.

5.2.2.2.2.1 Traffic Segmentation Overview

The upstream channels within the bonding group may have very different physical-layer characteristics. One channel may be 1 280 kbps with QPSK data regions and TDMA framing while another may be 5,12 Mbps with 64-QAM data regions and S-CDMA framing. The CMTS decides how to segment the bandwidth based on the bandwidth requested by the CM and the other traffic on the upstream channels. Figure 5-4 shows an example of four upstream TDMA channels with varying mini-slot sizes. Each row in the figure represents bandwidth across a single upstream channel. The vertical lines demarcate the mini-slot boundaries.

The letters and shadings in the figure represent the service flow to which the block of bandwidth has been allocated by the CMTS. Blocks E and D represent small grants to different flows supporting voice service. In this example, the CMTS chooses to grant A's request by using bandwidth on only Channels #1 and #2. Similarly the CMTS chooses to grant B's request by using only Channels #3 and #4. The CMTS chooses to grant C's request spread across all four upstream channels.




Figure 5-4: Segmentation Example

Each contiguous group of mini-slots assigned to the same service flow on the same channel in the figure becomes a segment. Thus the grant to service flow B consists of 2 segments and the grant to service flow C consists of 4 segments. Since the grant to service flow A on Channel #1 consists of two portions separated by the grant to service flow E, the overall grant to service flow A consists of 3 segments: two on Channel #1 and one on Channel #2. Each of these segments is treated like a legacy grant from the standpoint of physical layer overhead. Each segment will need a preamble at the beginning and, if TDMA transmission is used, guard time at the end. The physical layer properties of each segment are specified by the channel's physical parameters and the segment's burst parameters. The set of channels over which the CMTS may segment bandwidth for a given service flow is called the service flow's Upstream Bonding Group. The Upstream Bonding Group is used by the CMTS to know on which channels it may allocate grants to a service flow. The Upstream Bonding Group is also used by the CM to know on which channels it may send requests and on which channels it must look for grants for a given service flow.

5.2.2.2.2.2 Request/Grant Process

The request/grant mechanism for upstream channel bonding is an extension of the request/grant process used before DOCSIS 3.0. Prior to DOCSIS 3.0, CMs requested for individual packets or groups of packets and required a tight coupling between request and grants. The DOCSIS 3.0 introduces a packet streaming protocol called Continuous Concatenation and Fragmentation that allows a looser coupling between requests and grants and enables the CM to have multiple requests outstanding simultaneously. The CM requests bandwidth based on per-flow requirements such as queue-depth and QoS parameters. The CM may send bandwidth requests on any channel associated with the service flow and the CMTS may grant such a request on any combination of channels within the Upstream Bonding Group associated with the service flow.

When the CM transmits traffic for a service flow in a segment, it usually includes a segment header which contains a segment sequence number. The CMTS uses the segment sequence number to know the segment ordering for reassembling the service flow traffic stream.

5.2.2.3 Autonomous Load Balancing

Similar to DOCSIS 2.0 CMTS, the DOCSIS 3.0 CMTS supports autonomous load balancing of CMs. In DOCSIS 2.0 a mechanism was defined in which the CMTS could be configured with certain load balancing group information which would be used by the CMTS in order to balance load across a number of channels in the case where multiple channels reached a population of CMs. The load balancing group information described certain aspects of the plant topology that were necessary for the CMTS to perform the balancing operation.

In DOCSIS 3.0, the CMTS is configured with detailed plant topology information and the initialization procedure of the CM is designed such that the CMTS can locate (resolve) the CM's location in the plant topology. This is necessary for the support of channel bonding. Further, it is expected that most deployments will be configured such that multiple channels reach a population of CMs, so that the benefits of channel bonding can be realized. This leads to two important distinctions between the load balancing operations of DOCSIS 2.0 CMTSs and the load balancing operations of DOCSIS 3.0 CMTSs:

- 1) Balancing of pre-3.0 DOCSIS CMs: With a DOCSIS 3.0 CMTS, DOCSIS 2.0 and DOCSIS 1.1 CMs can be load balanced across the channels that physically reach those CMs. This would typically include the upstream channels and primary-capable downstream channels used by DOCSIS 3.0 CMs for channel bonding. The plant topology information used for channel bonding for DOCSIS 3.0 CMs is normally used for load balancing of pre-3.0 DOCSIS CMs. With a DOCSIS 2.0 CMTS, since complete plant topology information is not available and the CMTS does not attempt to resolve the topological location of CMs, certain topologies require the operator to configure (either via the CM configuration file or the CMTS directly) a priori information regarding a CM's expected plant location.
- 2) Balancing of DOCSIS 3.0 CMs: In certain deployments, there may be more channels that physically reach a set of DOCSIS 3.0 CMs than any individual CM can simultaneously receive. In this case, the DOCSIS 3.0 CMTS will balance the population of DOCSIS 3.0 CMs across the available channels by assigning each CM an appropriate subset of the channels upon which to operate. A DOCSIS 2.0 CMTS will treat DOCSIS 3.0 CMs just like DOCSIS 2.0 CMs, assigning a single upstream and a single downstream channel.

As in DOCSIS 2.0, the definition of "balanced" load is left to the CMTS vendor and the algorithm by which the CMTS attempts to achieve and maintain this balance is similarly left to the CMTS vendor.

5.2.3 Multicast Operation

DOCSIS 3.0 enhances support for IP Multicast with the addition of new features such as Source Specific Multicast [58], Quality of Service support for multicast traffic, IPv6 multicast and bonded multicast. These enhanced IP Multicast features enable cable operators to offer various IP Multicast-based multimedia services, such as Internet Protocol Television (IPTV), over the DOCSIS network. The following new features are added in DOCSIS 3.0 while maintaining backwards compatibility with the DOCSIS 2.0 multicast mode of operation:

- Forwarding of Source Specific Multicast (SSM) traffic for IGMPv3 [52] and MLDv2 [54] CPE devices.
- Support for bonded multicast traffic.
- Provisioning of Quality of Service (QoS) for multicast traffic.
- Support for IPv6 multicast traffic including Neighbor Discovery (ND), Router Solicitation (RS), etc.
- Explicit tracking of CPEs joined to a multicast group at the CMTS to aid load balancing, usage tracking, billing, etc.

DOCSIS 3.0 simplifies the operation of a Cable Modem (CM) by removing the IGMP snooping requirement of DOCSIS 1.1 and 2.0 (in some cases), instead of extending the use of IGMP snooping to support the above mentioned new features. The CM transparently forwards IGMP/MLD messages received from clients to the CMTS. A new CMTS-initiated layer-2 control mechanism is defined that configures the forwarding of downstream multicast packets to specific interfaces on the CM. The CMTS labels all multicast packets with a DSID (see clause 7.4). From the CMTS perspective, a DSID identifies a set of CMs intended to receive the same multicast packets. The CMTS communicates to a CM a DSID and associated group forwarding attributes, such as the set of CM interfaces to which these DSID-labeled multicast packets need to be forwarded. The same mechanism of DSID based filtering and forwarding is used for pre-registration as well as post-registration well-known IPv6 multicast traffic, such as Neighbor Discovery (ND) and Router Solicitation (RS). The CMTS can optionally encrypt multicast packets belonging to a particular multicast session using a Security Association (SA) communicated to a CM. Refer to clause 9.2, for further details.

QoS support for Multicast traffic is provided by leveraging already defined DOCSIS QoS constructs such as Service Flows and Classifiers. CMTS supports configuration of Group Classifiers that classify multicast packets intended to be received by a group of CMs on a downstream interface. Cable operators can statically provision Group Service Flows, clause 7.5.3, based on Service Class Names, clause 7.5.6, for multicast traffic. Refer to clause 7.5, QoS Support for further details.

5.2.4 Network and Higher Layer Protocols

At the Network Layer DOCSIS requires the use of Internet Protocol version 4 and version 6 for transporting management and data traffic across the HFC link between the CMTS and the CM.

As described above the CMTS could perform MAC Layer bridging or Network Layer routing of data traffic, while the CM only performs MAC layer bridging of data traffic. However both CMTS and CM are Network Layer and Transport Layer aware. Specifically, the CM and CMTS support classifying user traffic, based on Network Layer and Transport Layer criteria, for purposes of providing Quality of Service and packet filtering.

Additionally, DOCSIS requires use of the following Higher Layer Protocols for operation and management of the CM and CMTS:

- SNMP (Simple Network Management Protocol).
- TFTP (Trivial File Transfer Protocol), which is used by the modem for downloading operational software and configuration information.
- DHCP (Dynamic Host Configuration Protocol) v4 and v6, frameworks for passing configuration information to hosts on a TCP/IP network.

5.2.5 CM and CPE Provisioning and Management

5.2.5.1 Initialization, Provisioning and Management of CMs

During initialization, the CM goes through a number of steps before becoming fully operational on the DOCSIS network. The full initialization sequence is detailed in clause 10, but at a high level comprises four fundamental stages:

- 1) topology resolution and physical layer initialization;
- 2) authentication and encryption initialization;
- 3) IP initialization; and
- 4) registration (MAC layer initialization).

In the first stage, topology resolution and physical layer initialization, the CM acquires a single downstream channel (either via a stored last-known-good channel or by scanning the downstream channel map) and receives broadcast information from the CMTS that provides it with enough information to identify what set of downstream channels are available to it, as well as what upstream channels might be available. The CM then attempts to initialize the upstream physical layer by "ranging" on a selected upstream channel. Via a series of attempts and alternative channel selections, the CM succeeds in contacting the CMTS and completing the ranging process. At this point, the CMTS has located the CM in the plant topology (i.e. is aware of what downstream channels and upstream channels physically reach the CM) and has established two way communication via a single downstream/upstream channel pair. While this clause has referred to the first stage in terms of physical layer initialization, a provisional MAC layer initialization has been performed, with the full initialization of the MAC layer being deferred to the final stage.

The second stage, authentication and encryption initialization, involves the CM sending its X.509 digital certificate (including the CM's RSA public key) to the CMTS for validation. If the CM has sent a valid certificate, the CMTS will respond with a message that triggers the exchange of AES (or DES) encryption keys that are used to encrypt the upstream and downstream data transmissions from this point forward. This "Early Authentication and Encryption" can be disabled. If so, the CM will attempt authentication and encryption initialization after the registration stage. The details of the authentication and encryption initialization process are provided in [15].

In the third stage, IP initialization, the CM acquires an IP address in the Cable Operator address space, as well as the current time-of-day and a binary configuration file. DOCSIS 3.0 defines use of IP version 4 and IP version 6 and four provisioning modes: IPv4 Only, IPv6 Only, Alternate and Dual-stack. For IPv4 Only provisioning, the CM uses DHCPv4 to acquire an IPv4 address and operational related parameters. To facilitate compatibility with existing provisioning systems, this process is identical to the DOCSIS 2.0 CM provisioning process. For IPv6 Only provisioning, the CM uses DHCPv6 to acquire an IPv6 address and operational parameters. The CM uses the IPv6 address to obtain the current time-of-day and a configuration file. For Alternate Provisioning Mode (APM) the CM combines the first two provisioning modes, IPv6 Only and IPv4 Only, in sequential order, attempting IPv6 provisioning first and, if this fails, attempting IPv4 provisioning next. In the first three provisioning modes, IPv6 Only, IPv4 Only and APM, the CM operates with only one IP address type (v4 or v6) at any given time and thus these modes are called single-stack modes. For Dual-stack Provisioning Mode (DPM), the CM acquires both IPv6 and IPv4 addresses and parameters through DHCPv6 and DHCPv4 almost simultaneously, prioritizing the use of the IPv6 address for time-of-day and configuration file acquisition. In this mode, the CM makes both the IPv4 and the IPv6 addresses available for management.

The fourth stage, registration, involves a three-way handshake between the CM and the CMTS in which the CM passes certain contents of the configuration file to the CMTS, the CMTS validates the contents, reserves or activates MAC layer resources based on the service provisioning information that it received and communicates MAC layer identifiers back to the CM. Once the CM acknowledges receipt of the CMTS's response, the MAC layer initialization is complete.

After the CM completes initialization, it is a manageable network element in the operator's IP network. The CM supports SNMP (as mentioned above) and responds to queries directed to the IP (v4 or v6) address that it acquired during initialization. DOCSIS 3.0 also supports a dual-stack operational mode in which the CM is manageable via both IPv4 and IPv6 addresses simultaneously. This mode is initialized (i.e. the CM acquires a second IP address) after the CM is operational. This feature is also intended to help provide a streamlined migration from IPv4 to IPv6 in DOCSIS networks.

5.2.5.2 Initialization, Provisioning and Management of CPEs

DOCSIS assumes the use of DHCP for provisioning of CPE devices. To that end the CMTS supports a DHCP relay agent which allows the operator to associate a CPE IP Address request with the subscriber Cable Modem MAC Address. This feature is also used as the basis of a mechanism that prevents spoofing of IP Addresses.

DOCSIS 3.0 gives operator the option to provision CPE devices with an IPv4 or an IPv6 or both types of IP Addresses simultaneously.

5.2.6 Relationship to the Physical HFC Plant Topology

The basic connectivity principles for upstream and downstream connectivity between a CMTS and a CM are explained in annex M. Annex M explains how DOCSIS relates the HFC Plant Topology to CM Service Groups, MAC Domains and Bonding Groups.

5.2.6.1 RF Topology Configuration

CMTSs and CMs are interconnected by an RF combining and splitting network. A CMTS downstream channel is said to "reach" a CM when its downstream RF signal can be received by the CM. A CMTS upstream channel is said to "reach" a CM if the CMTS can receive the upstream transmission by that CM.

In most CMTS field deployments, the RF interconnection network is a Hybrid Fibre/Coaxial (HFC) network. An HFC network features a star wiring topology in which long distance fibres from a single head-end or hub location are distributed to fibre nodes throughout a geographic region. A fibre node usually terminates one or more downstream forward carrier paths from the head-end and originates one or more upstream reverse carrier path(s) to the head end. The fibre node connects the upstream and downstream signals from the fibre onto several coaxial cable segments (typically 2 to 4 segments). Multiple Cable Modems connect their single RF Port to the coax segment. The important topological feature of HFC networks is that all CMs connected to the same coax segment of a fibre node reach the same set of downstream and upstream channels on the CMTS(s) at the head-end.

The CMTS is configured with the physical topology of the plant. An operator configures the list of fibre nodes in the plant and configures which fibre nodes are reached by each downstream and upstream channel. A CMTS supports non-volatile configuration of a printable text name for each fibre node.

The operator also configures the set of MAC Domains in the CMTS and assigns each downstream and upstream channel to a MAC Domain. The CMTS automatically determines the MD-CM-SGs from the topology configuration of the operator.

Figure 5-5 depicts an example RF splitting/combining network to three fibre nodes. In this example, all channels are assumed to be configured to the same MAC Domain. Although the downstream connectivity is not typical, it has been chosen to demonstrate the flexibility of the topology configuration introduced with DOCSIS 3.0.

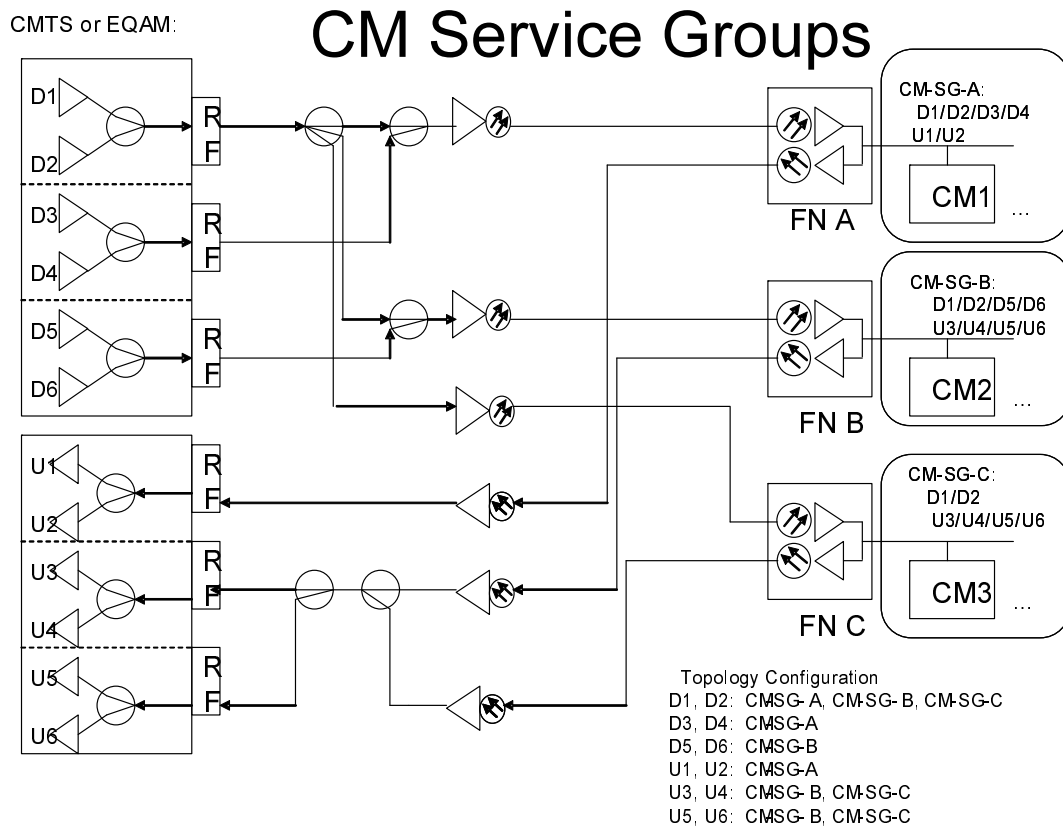


Figure 5-5: CM Topology Configuration Example

In figure 5-5, the CMTS implements six downstream channels organized as two Downstream RF channels per Downstream RF Port. The D1/D2 RF port is split three ways to reach to all three fibre nodes: nodes "FN-A", "FN-B" and "FN-C". The D3/D4 port reaches only the fibre node named "FN-A". The D5/D6 port reaches only fibre node named "FN-B".

The upstream from FN-A is connected to a single upstream RF port to which are attached receivers for separate upstream channels U1 and U2. For FN-B and FN-C, however, the signals from their upstream fibre are electrically combined and then split and connected to two CMTS RF ports. As a result, both fibre nodes "FN-B" and "FN-C" share the same set of upstream channels U3/U4/U5/U6.

The CMTS implements a "Node Configuration Table" management object with which an operator configures a textual name and number for each fibre node. The CMTS implements a "Topology Configuration Table" with which the operator configures which fibre nodes are reached by which downstream and upstream channels. The following tables represent the logical information of a Node Configuration Table and the Topology Configuration Table to describe the topology depicted in figure 5-5.

Table 5-1: Example Node Configuration Table

Node Number	Node Name
1	"FN-A"
2	"FN-B"
3	"FN-C"

Table 5-2: Example Topology Configuration Table

Node	Channel
1	D1
1	D2
1	D3
1	D4
1	U1
1	U2
2	D1
2	D2
2	D5
2	D6
2	U3
2	U4
2	U5
2	U6
3	D1
3	D2
3	U3
3	U4
3	U5
3	U6

For convenience, the "Channel" column of the example Topology Configuration Table above refers to the name from figure 5-6 to identify a channel. In actual practice, a channel is identified with an interface index with SNMP or with a (MAC Domain, channel ID) or other vendor-specific syntax to identify the channel with a CMTS vendor's command line interface.

5.2.6.2 Frequency Assignment

The topology database is configured at the CMTS to enable it to maintain frequency isolation for multiple channels reaching the same fibre node. During configuration, the CMTS will enforce that RF Channels reaching the same fibre node have different frequencies.

The CMTS uses the topology configuration to determine which channels can reach a CM for channel bonding, load balancing and multicast replication. Figure 5-6 shows a Frequency/Space diagram that depicts the reachability of downstream and upstream channels. This figure represents the same topology configuration as figure 5-5. In this figure, each vertical column on the left side of the figure (denoted by the labels DF_1 , DF_2 , DF_3 , DF_4) represents a downstream frequency, while each vertical column on the right side of the figure (denoted by the labels UF_1 , UF_2 , UF_3 , UF_4) represents an upstream frequency. Each rectangle (D1-D6 and U1-U6) represents a channel.

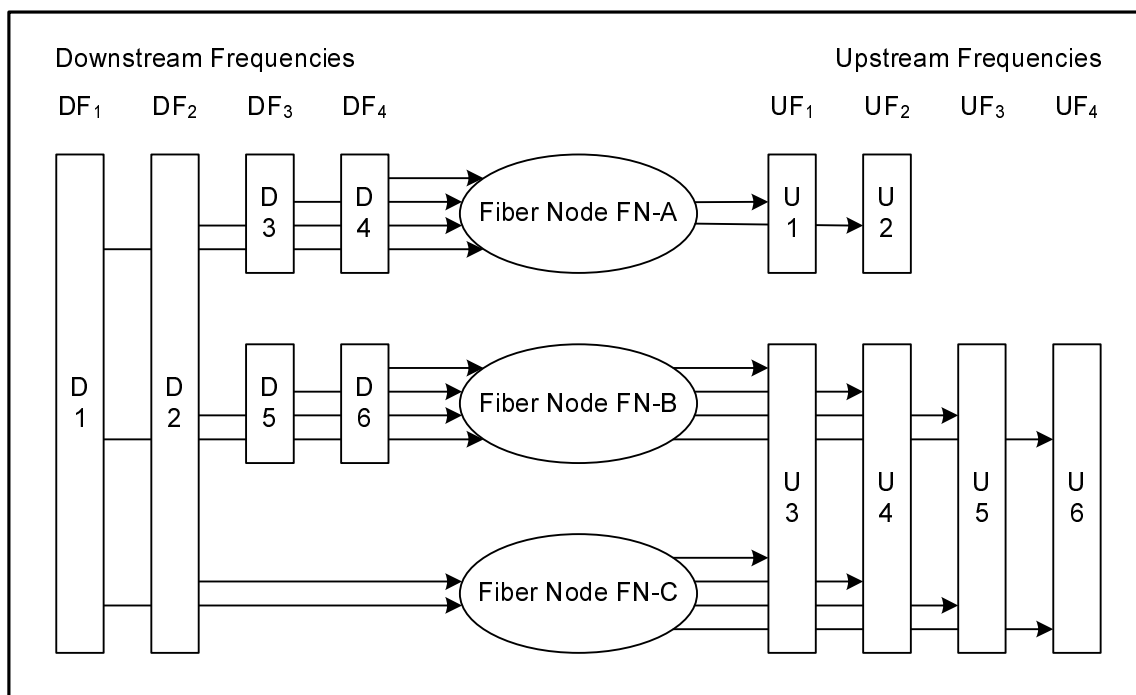


Figure 5-6: Frequency Space Diagram

5.2.7 Cable Modem Service Group (CM-SG)

A "Cable Modem Service Group" (CM-SG) is formally defined as the complete set of CMTS channels-both upstream and downstream-that reach a single cable modem. In an HFC deployment, all CMs reached by the same fibre node are reached by the same set of channels. Furthermore, in most HFC deployments, each fibre node has a different set of either upstream or downstream channels that reach it. Thus, a CM-SG usually corresponds to the channels reaching a single fibre node and the term "CM-SG" can generally be considered to be synonymous with "fibre node". In figure 5-6, for example, each of the fibre nodes: FN-A, FN-B and FN-C is a distinct CM-SG.

If two fibre nodes however, are reached by exactly the same set of downstream and upstream channels, then the CM-SG consisting of that set of channels is considered to contain both fibre nodes. An example of a CM-SG that contains two fibre nodes is depicted in the frequency/space below.




Figure 5-7: Multiple fibre nodes per CM-SG

A "Downstream Service Group" (DS-SG) is formally defined as the complete set of CMTS downstream channels that may be received by a single CM. A CM is reached by a single Downstream Service Group. A DS-SG represents a unique combination of DOCSIS Downstream RF Channels, each operating at a different center frequency. A DS-SG may be combined in the electrical domain and then be electrically and/or optically split to multiple fibre nodes. A DS-SG is a set of channels defined by the topology configuration of the CMTS and is independent of the MAC Domain configuration.

An "Upstream Service Group" (US-SG) is formally defined as the complete set of upstream channels in a CMTS that may receive the transmissions of a single CM. A US-SG is a physical-layer concept; it is defined only by the physical combining of the upstream RF transmission from CMs. If the upstream fibre signals from different fibre nodes are not combined, each fibre node usually corresponds to a single US-SG.

NOTE: A CM-SG, DS-SG and US-SG are completely defined by the topology configuration of CMTS channels and fibre nodes reached by them. These terms are independent of the assignment of channels to MAC Domains.

5.2.7.1 MAC Domain Channel Assignment

An operator configures each upstream and downstream channel of a CMTS into a MAC Domain. In a frequency/space diagram, a MAC Domain can be represented by a "barbell" that encloses the downstream channels of the MAC Domain on one side and the upstream channels of the MAC Domain on the other side.

Figure 5-8 shows a typical topology with three fibre nodes and two MAC Domains.

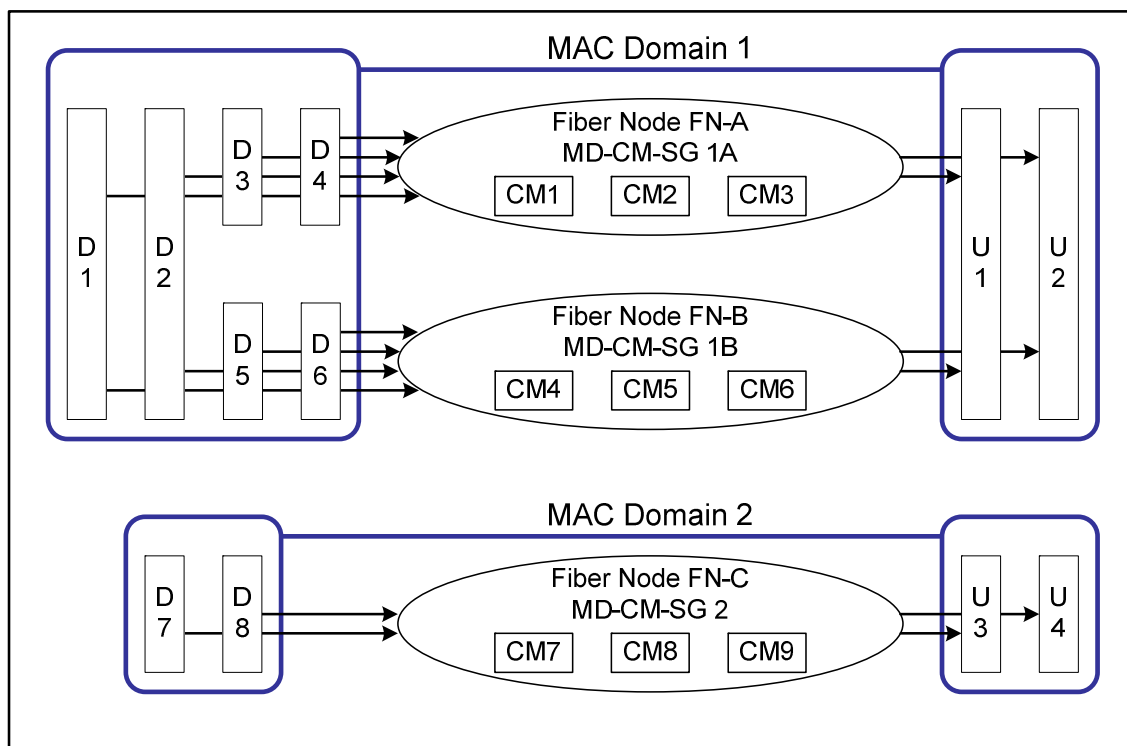


Figure 5-8: Example MAC Domain Channel Assignment

In this example, downstream channels D1 and D2 reach both FN-A and FN-B, while downstream channels D3/D4 reach only FN-A and D5/D6 reach only FN-B. MAC Domain 1, which includes D1/D2, reaches both fibre node FN-A and FN-B. MAC Domain 1 consists of all of the channels D1/D2/D3/D4/D5/D6/U1/U2. Notice that the CMs in FN-A can reach only the channels D1/D2/D3/D4/U1/U2, while the CMs in FN-B reach a different set of channels D1/D2/D5/D6/U1/U2.

A "MAC Domain CM Service Group" (MD-CM-SG) is the set of downstream and upstream channels from the same MAC Domain, all of which reach a single CM. An MD-CM-SG corresponds to a general load balancing group because it forms the set of channels among which a non-bonding CM can be moved while remaining registered in the same MAC Domain. For bonding CMs, an MD-CM-SG represents the set of channels among which traffic on bonded service flows can be scheduled while the CM remains registered to the same MAC Domain.

The channels configured for MAC Domain 2 are D7/D8/U3/U4. These channels reach only fibre node FN-C. MAC Domain 2 has only one MD-CM-SG, with channels D7/D8/U3/U4.

Because a MAC Domain defines a separate address space for many DOCSIS protocol elements (e.g. DSIDs, SAIDs, etc.), an operator should define separate MAC Domains that serve disjoint subsets of fibre nodes rather than a single MAC Domain for all fibre nodes.

A CMTS implementation may restrict the configuration of the downstream channels and upstream channels in the same MAC Domain.

DOCSIS 3.0 introduces a mechanism whereby the CMTS determines the MD-CM-SG of a CM when it registers (see clause 10). If each MD-CM-SG corresponds to a single fibre node, the CMTS can thereby determine the fibre node that reaches each registered CM. An MD-CM-SG always contains at least one fibre node.

5.2.7.2 Multiple MAC Domains per Fibre Node

For simplicity, it is recommended that all DOCSIS channels from a CMTS reaching a fibre node be configured into the same MAC Domain. It may be desired, however, to define separate sets of downstream and upstream channels that reach the same fibre node into different MAC Domains in order to provide separate services. For example, business customers or set-top-box CMs may be desired to have entirely separate service from residential high-speed-data CMs and may be configured onto separate MAC Domains.

Figure 5-9 shows an example of two MAC Domains implemented on the different downstream and upstream channels that reach the same set of fibre nodes.

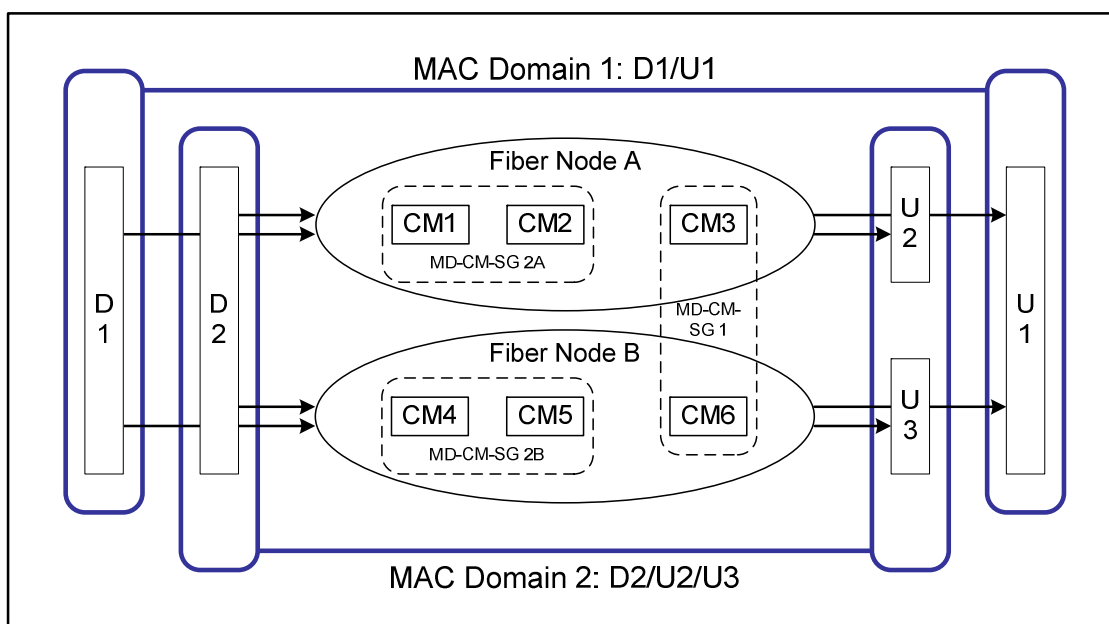


Figure 5-9: Multiple MAC Domains per Fibre Node

In the above example, the topology is such that downstream channels D1 and D2 reach both FN-A and FN-B. Upstream channel U1 is reached by FN-A and FN-B, but U2 is reached only by FN-A and U3 only by FN-B. The operator desires that set-top boxes in both fibre nodes use the "high split" (2 FNs per channel) channels D1 and U1 and for residential CMs in both fibre nodes to use the "low split" (1 FN per channel) channels D2, D3, U2 and U3.

The operator configures MAC Domain 1 to contain channels D1 and U1 and for MAC Domain 2 to contain channels D2, U2 and U3. This causes the formation of three "MAC Domain CM Service Groups". MD-CM-SG 1 consists of the channels D1/U1, i.e. the channels of MAC Domain 1, which reaches two fibre nodes. Note that when a set-top-box registers on MAC Domain 1, the CMTS can not tell whether the CM is physically connected in FN-A or FN-B. MD-CM-SG 2A consists of channels D2/U2, i.e. the channels in MAC Domain 2 that reach fibre node A. MD-CM-SG 2B consists of channels D2/U3, i.e. the channels in MAC Domain 2 that reach fibre node B.

5.2.7.3 MAC Domain Downstream and Upstream Service Groups

The term "MAC Domain Downstream Service Group" (MD-DS-SG) refers to the set of downstream channels from the same MAC Domain that reaches a fibre node. In many cases, an operator will configure all downstream channels reaching a fibre node to the same MAC Domain, in which case an MD-DS-SG corresponds to a DS-SG from the topology configuration.

In general, an MD-DS-SG may contain downstream channels that are shared by multiple MD-CM-SGs, each of which has a different upstream channel. In the example shown in figure 5-9, MAC Domain 2 has only a single MD-DS-SG (containing D2), which contains the downstream channels of two MD-CM-SGs.

The term "MAC Domain Upstream Service Group" (MD-US-SG) refers to the set of upstream channels from the same MAC Domain that is reached by a single CM. In the common case where all of the upstream channels reached by a fibre node are configured in the same MAC Domain, an MD-US-SG corresponds to a US-SG defined by the topology configuration.

In general, an MD-US-SG may contain upstream channels shared by multiple MD-CM-SGs, each of which has a different set of downstream channels. In the example shown in figure 5-8, MAC Domain 1 has a single MD-US-SG (containing U1/U2) which contains the upstream channels of two MD-CM-SGs.

5.2.7.4 Channel Bonding Topology Considerations

A "Provisioned" Bonding Group is a configured set of downstream or upstream channels on the same MAC Domain that reach at least one fibre node in common. Figure 5-10 takes the Service Groups and MAC Domains defined in earlier figures and overlays a variety of provisioned Downstream Bonding Groups (DBG) and Upstream Bonding Groups (UBG). In addition to provisioned bonding groups, a CMTS may dynamically create downstream or upstream bonding groups.

Because a single CM must be able to reach all channels of a bonding group, a CMTS SHOULD restrict configuration of provisioned bonding groups such that all channels reach at least one fibre node in common.

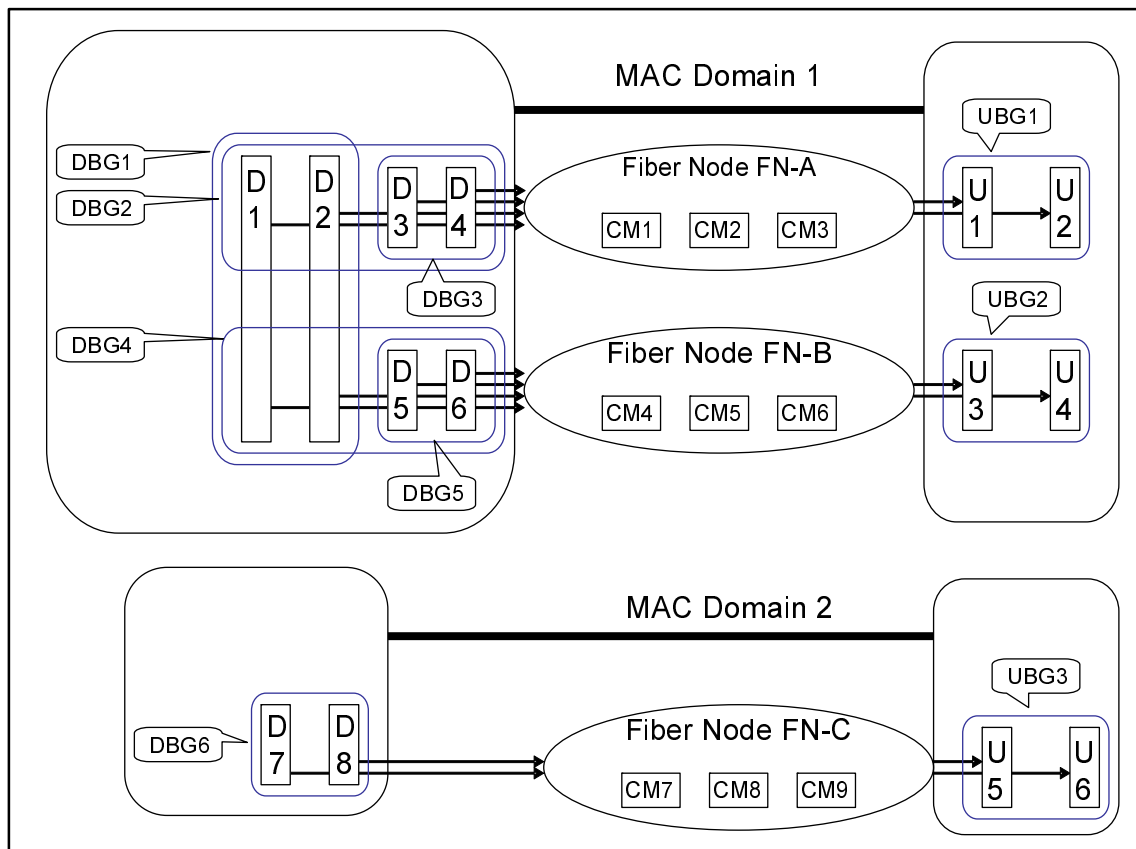


Figure 5-10: Bonding Group Example

The Downstream Bonding Groups depicted include DBG1{D1, D2}, DBG2{D1, D2, D3, D4}, DBG3{D3, D4}, DBG4{D1, D2, D5, D6}, DBG5{D5, D6} and DBG6{D7, D8}. The Upstream Bonded Channel Sets depicted include UBG1{U1, U2}, UBG2{U3, U4} and UBG3{U5, U6}.

A CMTS may restrict the set of channels assigned to a Bonding Group based on vendor implementation. For example, a CMTS may require that bonded RF channels reside on RF ports of the same line card or even on the same RF Port.

For downstream multicast forwarding, an important concept is a "Downstream Channel Set" (DCS). A DCS is either a single downstream channel or a downstream bonding group. A downstream multicast session is said to be replicated onto a DCS, i.e. it is transmitted either on a single downstream channel or transmitted on the multiple channels of a downstream bonding group. In the example of figure 5-10, there are a total of 14 DCSs: eight individual downstream channels and six downstream bonding groups. A downstream multicast session can be replicated on any or all of the 14 DCSs of the example topology.

5.2.8 CMTS Downstream Service Model Example

The model for downstream forwarding with bonding groups is an extension of the MAC service model for the CMTS. The model remains that downstream bonded service is offered by MAC Domains and that the "CMTS Forwarder" is responsible for forwarding packets from a Network Side Interface (NSI) to the MAC Service Access Point (MSAP) of one or more MAC Domains.

An example DOCSIS Downstream Service Model is depicted in figure 5-11.

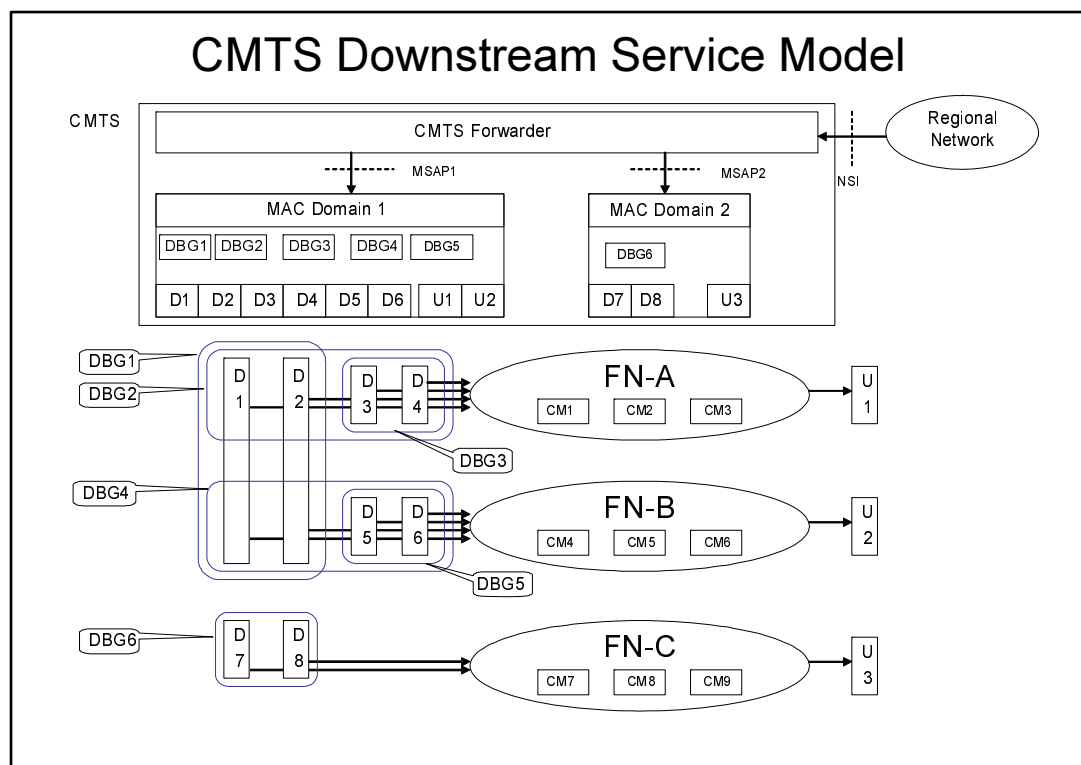


Figure 5-11: CMTS Downstream Service Model

This is a conceptual model that describes operations of internal functions in the CMTS and CM that result in external behavior on the interfaces of the products. It is intended to clarify interpretation of normative requirements of external behavior on RF interface and OSSI interface. This model is not intended to describe or restrict the implementation of these functions in actual products. A product may have internal implementation in any manner consistent with the normative requirements of external behavior.

In the model, a "CMTS Forwarder" subcomponent is modeled as having already constructed a layer 2 Ethernet Packet PDU for downstream transmission and delivered it to a MAC Domain's MAC Service Access Point for downstream forwarding service. Furthermore, the CMTS Forwarder has determined whether the packet is to be forwarded to a single CM or to multiple CMs and if to a single CM, the service request includes an internal identifier of the CM.

Operation of IP layer 3 forwarding, as well as IGMP and IP multicast forwarding, is modeled as an operation of the CMTS Forwarder, not of the MAC Domain.

The semantics of the MAC Domain's MAC Layer service primitives are different for unicast traffic intended to an individual CM and multicast traffic intended for delivery to a group of CMs. The MAC service level primitives are described in annex L. For traffic to an individual CM, the CMTS Forwarder is considered to identify the CM when it provides the packet to the MAC Domain. For traffic to a group of CMs, the CMTS Forwarder is considered to identify the Downstream Channel Set on which the MAC Domain is to transmit the packet.

The CMTS Forwarder is responsible for determining which MAC Domains and which Downstream Channel Sets reach the desired hosts of a multicast downstream packet. Desired hosts include embedded CM hosts and CPE hosts reachable through a CM's CMCI interface. The CMTS Forwarder is responsible for determining how downstream multicast packets are replicated to the multiple downstream channel sets of each MAC Domain.

6 Media Access Control Specification

6.1 Introduction

6.1.1 Overview

This clause describes version 3.0 of the DOCSIS MAC protocol. Some of the highlights of the DOCSIS MAC protocol 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.
- Logical combining of multiple channels for increased throughput (channel bonding).

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 logical upstream channels and one or more downstream channels. The CMTS **MUST** discard any packets received that have a source MAC address that is not a unicast MAC address. The upstream channels may be any combination of DOCSIS 1.x, 2.0 or 3.0 formats. A single upstream channel may transport DOCSIS 1.x, 2.0 and 3.0 bursts.

6.1.2.2 MAC Service Access Point

A MAC Service Access Point (MSAP) is an attachment to a MAC-sublayer domain (refer to clause 9.1.1).

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, Service Flows 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 assigns 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 clause 11.2).

For example, in a basic CM implementation, two Service Flows (one upstream, one downstream) could be used 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 are referred to as the upstream and downstream Primary Service Flows. The Primary Service Flows needs to always be provisioned to allow the CM to request and to send the largest possible unconcatenated MAC frame (refer to clause 6.2.2).

The CM and CMTS **MUST** immediately activate the Primary Service Flows at registration time. The CM and CMTS **MUST** always use the Ranging SID(s) for periodic ranging after registration when in Multiple Transmit Channel Mode and the Primary SID when not in Multiple Transmit Channel Mode. The Primary Service Flows may be used for traffic. All unicast Service Flows use the security association defined for the Primary Service Flow (refer to [15]).

The CMTS **MUST** ensure that all Service Flow IDs are unique within a single MAC-sublayer domain. An active/admitted service flow maps to one or more SIDs. SIDs are unique per logical upstream channel. 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).

Unicast flows on different logical upstreams that are attached to a single MAC-sublayer domain **MAY** be assigned the same SID by the CMTS, as long as the SFIDs are unique.

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. The usage code values are defined in table 6-28 and their allowed use is defined in clause 6.4. The binding of these values to physical-layer parameters is defined in table 6-26.

6.1.2.4.1 TDMA mode

For DOCSIS 1.x channels, a mini-slot is a power-of-two multiple of 6,25 microsecond increments, limited to 2, 4, 8, 16, 32, 64 or 128 times 6,25 microseconds. For DOCSIS 2.0 and 3.0 TDMA, a mini-slot is a power-of-two multiple of 6,25 microsecond increments limited to 1, 2, 4, 8, 16, 32, 64 or 128 times 6,25 microseconds. The relationship between mini-slots, bytes and time ticks is described further in clause 7.1.4.1.

6.1.2.4.2 S-CDMA mode

For DOCSIS 2.0 and 3.0 S-CDMA channels, a mini-slot is not restricted to be a power-of-two multiple of 6,25 microsecond increments. Instead a mini-slot is a unit of capacity that is dependent on the modulation rate, number of spreading codes and number of spreading intervals configured for the upstream channel. While the channel may be configured such that the time duration of a mini-slot is a power-of-two multiple of 6,25 microsecond increments, there is no special significance to 6,25 microsecond time ticks for S-CDMA channels. The relationship between mini-slots and S-CDMA framing is described further in [12]. The relationship between mini-slots, bytes and time ticks is described further in clause 7.1.4.2.

6.1.2.5 MAC Frame

A MAC 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 48-bit source and destination MAC addresses, data and a CRC. In special cases, the MAC Header may encapsulate multiple MAC frames (see clause 6.2.5.6) into a single MAC frame. The MAC layer definition of a frame is different from any physical layer or MPEG layer definition of a frame.

6.1.2.6 Logical Upstream Channels

The MAC layer deals with logical upstreams. A logical upstream is identified with an upstream channel ID which is unique within the MSAP. A logical upstream consists of a contiguous stream of mini-slots which are described by UCD messages and allocated by MAP messages associated with a channel ID. A CM operating with Multiple Transmit Channel Mode disabled can only register to operate on a single logical upstream channel. A CM in Multiple Transmit Channel Mode of operation can register to operate on one or more logical upstream channels.

There are four distinct types of logical upstreams:

- 1) Type 1: DOCSIS 1.x upstreams that support no DOCSIS 2.0 TDMA features.
- 2) Type 2: Mixed upstreams that support DOCSIS 1.x and DOCSIS 2.0 TDMA bursts.
- 3) Type 3: DOCSIS 2.0 upstreams, which cannot support DOCSIS 1.x CMs and include the following two subtypes:
 - a) Type 3A: DOCSIS 2.0 TDMA upstreams.
 - b) Type 3S: DOCSIS 2.0 S-CDMA upstreams.
- 4) Type 4: DOCSIS 3.0 upstreams, some of which cannot support Pre-3.0 DOCSIS CMs and include the following four subtypes:
 - a) Type 4A: The TDMA upstream is described by Type 29 UCDs for 2.0 CMs using IUCs 9, 10 and 11 for data grants and by Type 35 UCDs for 3.0 CMs using IUCs 5, 6, 9, 10 and 11 for data grants.
 - b) Type 4S: The S-CDMA upstream is described by Type 29 UCDs for 2.0 CMs using IUCs 9, 10 and 11 for data grants and by Type 35 UCDs for 3.0 CMs using IUCs 5, 6, 9, 10 and 11 for data grants.
 - c) Type 4AR: The DOCSIS 3.0 TDMA upstream is described by Type 35 UCDs for 3.0 CMs using IUCs 5, 6, 9, 10 and 11 for data grants. These channels are restricted to only DOCSIS 3.0 CMs.
 - d) Type 4SR: The DOCSIS 3.0 S-CDMA only upstream is described by Type 35 UCDs for 3.0 CMs using IUCs 5, 6, 9, 10 and 11 for data grants. These channels are restricted to only DOCSIS 3.0 CMs and have the option of using Selectable Active Codes Mode 2 and Code Hopping Mode 2 (see clause 6.4.3, Upstream Channel Descriptor (UCD) and [12]).

All valid logical upstreams fall into one of these 8 categories: Type 1, Type 2, Type 3A, Type 3S, Type 4A, Type 4S, Type 4AR or Type 4SR.

A CM operating in Multiple Transmit Channel Mode can use any of these logical channel types. However, when selecting the first upstream channel to use, the CM preferentially makes the selection based on the requirements in clause 10.2.3.

DOCSIS 2.0 introduced the possibility for multiple logical upstreams to share the same spectrum. When this occurs the logical upstreams sharing the same spectrum are time domain multiplexed and only one is active at any time, with the exception that it is possible for the Broadcast Initial Maintenance regions to be simultaneous. When a logical upstream channel is inactive, its mini-slots are allocated to the NULL SID by its associated MAP messages. Having multiple logical upstreams that share the same spectrum is the only way to have modems operating in S-CDMA mode share the same upstream spectrum with modems not operating in S-CDMA mode. Also, having multiple logical upstreams that share the same spectrum is the only way to have modems operating in S-CDMA mode with Selectable Active Codes Mode 2 enabled share the same upstream spectrum with other modems operating in S-CDMA mode without Selectable Active Codes Mode 2 enabled. Thus, it is possible in DOCSIS 3.0 to have three logical channels in the same upstream spectrum: one for a DOCSIS 3.0 only operation with Selectable Active Codes Mode 2 enabled, one for DOCSIS 3.0 only operation with Selectable Active Codes Mode 2 disabled and one for modems not operating in S-CDMA mode.

The CMTS MUST support the first four categories of logical upstream channel types individually. If the CMTS supports Selectable Active Codes Mode 2 ([12]) or supports assignment of advanced burst profiles for data associated with IUCs 5, 6, 9, 10 or 11, the CMTS MUST support the type 4 logical channel.

The CMTS MUST support the following combinations of logical upstream channels sharing the same upstream spectrum:

- One Type 1 logical channel plus one Type 3S logical channel with the same modulation rate on both logical channels.
- One Type 2 logical channel plus one Type 3S logical channel with the same modulation rate on both logical channels.
- One Type 3A logical channel plus one Type 3S logical channel with the same modulation rate on both logical channels.

The CMTS MAY support a Type 4 logical channel individually. The CMTS MAY support other combinations of logical channels sharing the same upstream spectrum, including Type 4 channels and combinations of logical channels with different modulation rates.

6.1.2.6.1 Type 3 Logical Upstreams

Type 3 Logical Upstreams have operational parameters in their associated UCD messages that prevent the operation of DOCSIS 1.x CMs. See clause 6.4.3 for a detailed description of which parameter values make a channel a Type 3A or 3S Upstream. The UCD messages for Type 3 Logical Upstreams use a different MAC management message type (see clause 6.4.1) than do UCD messages for channels that can support 1.x CMs. This prevents 1.x CMs from attempting to use Type 3 Upstreams or from being confused by UCD messages for those channels. A logical upstream is a Type 3A upstream if and only if it is described by a Type 29 UCD with version 3, does not contain burst profiles for IUC 5 and 6 and is a DOCSIS 2.0 TDMA upstream. A logical upstream is a Type 3S upstream if and only if it is described by a Type 29 UCD with version 3, does not contain burst profiles for IUC 5 and 6 and is an S-CDMA upstream without Selectable Active Codes Mode 2.

6.1.2.6.2 Type 4 Logical Upstreams

Type 4 Logical Upstreams are identified by UCD Type 35 and may additionally have UCD Type 29. The presence of UCD Type 29 allows use of these logical upstream channels by DOCSIS 2.0 CMs. If the UCD Type 29 is not present, the channel is restricted to use by DOCSIS 3.0 CMs only.

This channel type allows the operator to define burst profiles for five data IUCs (5, 6, 9, 10 and 11) for use by DOCSIS 3.0 CMs. The CMTS is free to select, using proprietary criteria, the most appropriate data IUC for each data burst for 3.0 CMs operating in Multiple Transmit Channel Mode. If UCD Type 29 is present, the operator should configure IUCs 9 and 10 to be appropriate for short and long data bursts for DOCSIS 2.0 CMs.

Additionally, Type 4SR logical upstreams allow the use of Selectable Active Codes Mode 2 and Code Hopping Mode 2 (see clause 6.4.3 and [12]).

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 the present document. These fields will not be interpreted or used in any manner by this version (3.0) of the MAC protocol.

The CMTS MUST transmit all Reserved or "for future use" fields as zero. The CM MUST silently ignore all Reserved or "for future use" fields.

The CM MUST transmit all Reserved or "for future use" fields as zero. The CMTS MUST silently ignore all Reserved or "for future use" fields.

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.

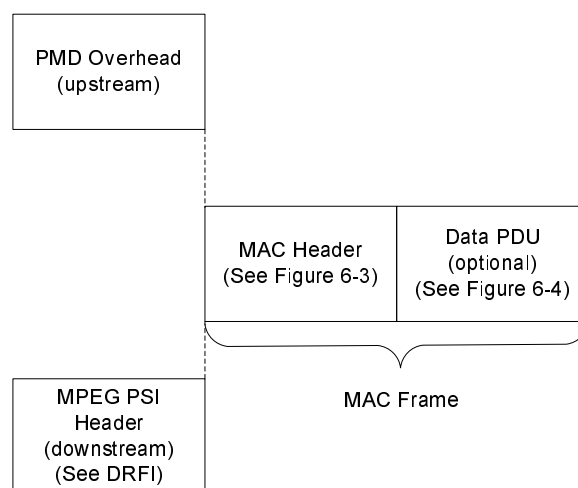


Figure 6-1: Generic MAC Frame Format

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 clause of [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 clause of the present document (refer to clause 7.2.1).

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.

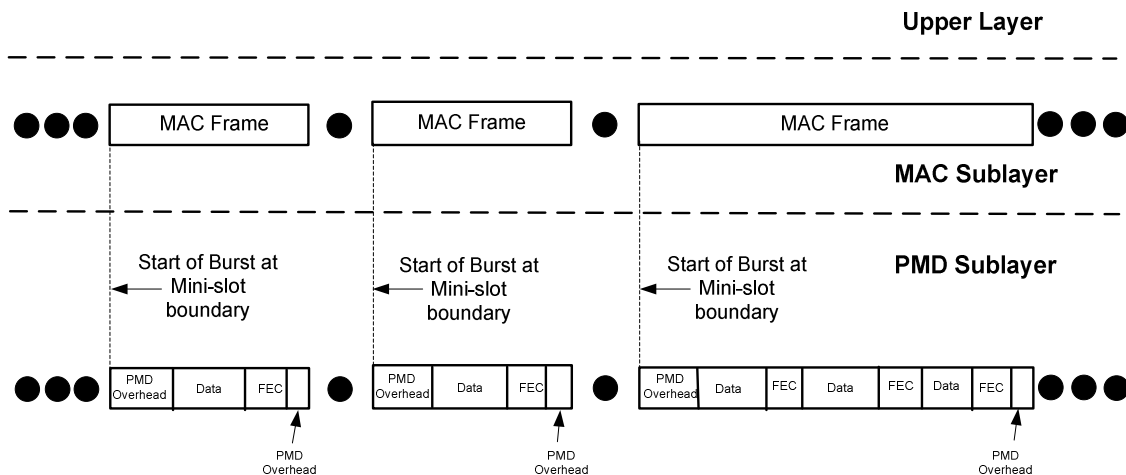


Figure 6-2: Upstream MAC/PMD Convergence

The layering of MAC frames over MPEG in the downstream channel is described in [4].

NOTE: The CMTS PHY ensures that, for a given channel, the CMTS MAC receives upstream MAC frames in the same order the CM mapped the MAC frames onto mini-slots. That is to say that if MAC frame X begins in mini-slot n and MAC frame Y begins in mini-slot $n+m$, then the CMTS MAC will receive X before it receives Y. This is true even when, as is possible with S-CDMA, mini-slots n and $n+m$ are actually simultaneously transmitted within the PHY layer.

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 [18]. This is often called bit-little-endian order.

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.

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 specification uses the following textual conventions:

- When tables describe bit fields within an octet, 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.
- When figures depict bit positions within an octet, the most significant bits are leftmost in the figure. For example, see the locations of the FC_TYPE and EHDR_ON bits in figure 6-3.
- When bit-strings are presented in text, the most significant bit is leftmost in the string.
- Unless explicitly indicated otherwise, when bits are enumerated in a bit-field, the least significant bit of the bit-field is bit # 0. The exceptions are certain fields that utilize the BITS Encoding convention.
- When message formats are presented in figures, the message octets are shown in the order in which they are transmitted on the wire, beginning with the field in the upper left and reading left-to-right, one row at a time. For example, in figure 6-13, the FC byte is transmitted first, followed by the MAC PARM and LEN fields. As mentioned above, the LEN field is transmitted with most-significant octet first and each octet is transmitted with least-significant bit first.

6.2.1.3.1 Representing Negative Numbers

Signed integer values **MUST** be transmitted and received by the CM and CMTS in two's complement format.

6.2.1.3.2 Type-Length-Value Fields

Many MAC messages incorporate Type-Length-Value (TLV) fields. Except for the cases of Primary Service Flow selection and MIC calculation among the TLVs encoded in a CM Configuration File, TLV fields are unordered lists of TLV-tuples. Some TLVs are nested (see annex C). The CM or CMTS **MUST** set all TLV Length fields, except for EH_LEN (see clause 6.2.5), to 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 not treat the event as an error condition.

6.2.1.4 MAC Header Format

The CM or CMTS **MUST** use the MAC Header format as shown in figure 6-3.

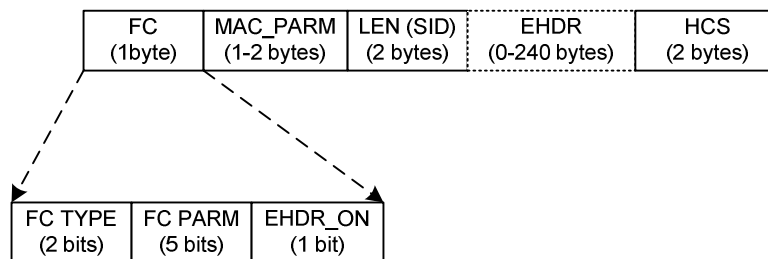


Figure 6-3: MAC Header Format

The CM **MUST** comply with table 6-1 for all MAC Headers. The CMTS **MUST** comply with table 6-1 for all MAC Headers. The Frame Control (FC) field is the first byte and uniquely identifies the rest of the contents within the MAC Header. The FC field is followed by 3 bytes of MAC control; an optional Extended Header field (EHDR); plus a Header Check Sequence (HCS) to ensure the integrity of the MAC Header.

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 for all headers except for the Queue-Depth based request header in which this field is 16 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 byte to 240 bytes
HCS	MAC Header Check Sequence	2 bytes
	Length of a MAC Header	6 bytes + EHDR

FC Field: The FC field is broken down into the FC_TYPE sub-field, FC_PARM sub-field and an EHDR_ON indication flag. The CM **MUST** comply with the FC field in table 6-2. The CMTS **MUST** comply with the FC field in table 6-2 for the FC field.

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: Isolation Packet 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 includes the two MSBs of the FC field. These bits MUST always be interpreted by CMs and CMTSs 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 with packet PDU Isolation from Pre-3.0 DOCSIS cable modems; or a MAC Header used for specific MAC control purposes. These types are spelled out in more detail in the remainder of this clause.

The five bits following the FC_TYPE sub-field is the FC_PARM sub-field. The use of these bits is 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.

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

MAC_PARM: 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 by the CM and CMTS as the Extended Header length (ELEN). The EHDR field may vary from 0 byte 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 clause 6.2.5.6). If this is a Request MAC Header (REQ) (see clause 6.2.4.3), 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.

LEN (SID): 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.

EHDR: 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.

HCS: The HCS field is a 16-bit CRC that ensures the integrity of the MAC Header, even in a collision environment. The CM or CMTS MUST include the entire MAC Header, starting with the FC field and including any EHDR field that may be present for HCS field coverage. The HCS is calculated using CRC-CCITT ($x^{16} + x^{12} + x^5 + 1$) as defined in [21].

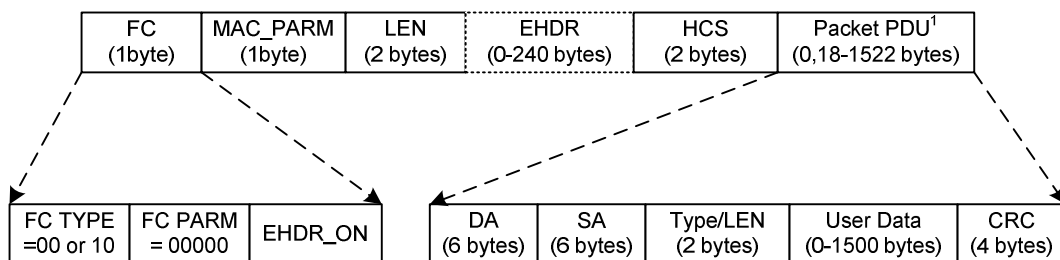
6.2.1.5 Data PDU

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, an Isolation Packet Data PDU and a MAC-Specific Frame. 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 Packet PDU and Isolation Packet PDU

The CM or CMTS MAC sublayer MUST support both, a variable-length Ethernet/[18]-type Packet Data PDU MAC Frame and a variable-length Ethernet/[18]-type Isolation Packet Data PDU MAC Frame. The Isolation Packet Data PDU MAC Frame is used to prevent certain downstream packets from being received and forwarded by Pre-3.0 DOCSIS cable modems, as described in clause 9.2.2.1. Both the Packet PDU and the Isolation Packet PDU can be used to send packets of any type (unicast, multicast and broadcast). With the exception of packets which have been subject to Payload Header suppression, the Packet PDU MUST be passed across the network in its entirety, including its original CRC. In the case where Payload Header Suppression has been applied to the Packet PDU, all bytes except those suppressed MUST be passed across the network by the CM and CMTS and the CRC covers only those bytes actually transmitted (refer to clause 6.2.5.4.1). A unique Packet MAC Header is appended to the beginning. The CM MUST comply with figure 6-4 and table 6-3 for Packet PDUs and Isolation Packet PDUs. The CMTS MUST comply with figure 6-4 and table 6-3 for Packet PDUs and Isolation Packet PDUs.



¹ Packet PDU length is limited to 1518 bytes in the absence of VLAN tagging. When PHS is applied, it is possible for the Packet PDU length to be less than 18 bytes.

Figure 6-4: Packet PDU or Isolation Packet PDU MAC Frame Format

Table 6-3: Packet PDU or Isolation Packet PDU MAC Frame Format

Field	Usage	Size
FC	FC_TYPE = 00; Packet PDU MAC Header FC_TYPE = 10; Isolation Packet PDU 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 to 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 [18] Length Field User Data (variable length, 0 byte to 1 500 bytes) CRC - 32-bit CRC over packet PDU (as defined in Ethernet/[18])	n bytes
	Length of Packet PDU or Isolation Packet PDU MAC frame	6 + x + n bytes

Under certain circumstances 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, e.g. a 5-byte Downstream Service Extended Header containing the current Sequence Number for a particular DSID (also known as a "null packet") or a Service Flow Extended Header containing the number of active grants for a UGS-AD service flow. This could also happen as a result of PHS in the upstream direction (see PHS clause 6.2.5.4.1).

Such a frame will have the length field in the MAC header set to the length of the extended header and will have no packet data and therefore no CRC.

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 CMs and CMTSs compliant with this version (3.0) of the specification. Compliant version 3.0 CM and CMTS implementations MUST use the length field to skip over the ATM PDU.

6.2.4 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 adjustment, requesting bandwidth, fragmentation and concatenating multiple MAC frames.

Table 6-4 describes FC_PARM usage within the MAC Specific Header.

Table 6-4: MAC-Specific Headers and Frames

FC_PARM	Header/Frame Type
00000	Timing Header
00001	MAC Management Header
00010	Request Frame
00011	Fragmentation Header
00100	Queue Depth-based Request Frame
11100	Concatenation Header

6.2.4.1 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 by the CMTS to transport the Global Timing Reference to which all cable modems synchronize. In the upstream, this MAC Header MUST be used by the CM 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 CM MUST comply with figure 6-5 and table 6-5 for Timing Headers. The CMTS MUST comply with figure 6-5 and table 6-5 for Timing Headers.

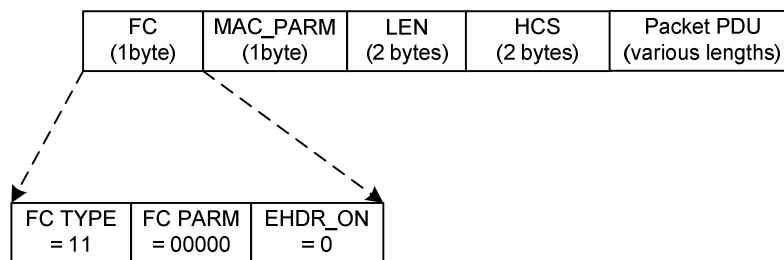


Figure 6-5: Timing MAC Header

Table 6-5: 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 byte
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

6.2.4.2 MAC Management Header

A specific MAC Header is identified to help support the MAC management messages required. This MAC Header MUST be used by CMs and CMTSs to transport all MAC management messages (refer to clause 6.4). The CM MUST comply with figure 6-6 and table 6-6 for MAC Management Headers. The CMTS MUST comply with figure 6-6 and table 6-6 for MAC Management Headers.




Figure 6-6: Management MAC Header

Table 6-6: 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 to 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

6.2.4.3 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. The CM MUST NOT include any Data PDUs following the Request Frame. The CM MUST comply with figure 6-7 and table 6-7 for Request Frames.




Figure 6-7: Request Frame Format

Table 6-7: 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 mini-slots requested	8 bits
SID	Service ID used for requesting bandwidth. For valid SID ranges, see clause 7.2.1.2	16 bits
EHDR	Extended MAC Header not allowed	0 byte
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 CM MUST replace the LEN field with a SID. The SID uniquely identifies a particular Service Flow within a given CM.

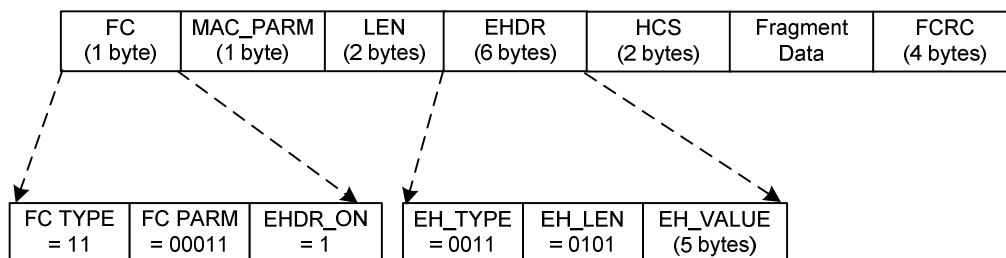
The CM MUST specify the bandwidth request, REQ, in mini-slots. The CM MUST indicate the current total amount of bandwidth requested for this service queue including appropriate allowance for the PHY overhead in the MAC_PARM field.

The Request Frame is for Pre-3.0 DOCSIS support and MUST NOT be used by CMs operating in Multiple Transmit Channel Mode. CMs operating in Multiple Transmit Channel Mode MUST use queue depth based requests as defined in clause 6.2.4.5.

6.2.4.4 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, Fragmentation is only applicable in the upstream. The CM MUST comply with figure 6-8 and table 6-8 for Fragmentation MAC Headers.

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

**Figure 6-8: Fragmentation MAC Header Format****Table 6-8: 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 clause 6.2.5.3	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/[18])	4 bytes
	Length of a MAC Fragment Frame	16 + n bytes

The Fragmentation MAC Frame is for Pre-3.0 DOCSIS support and MUST NOT be used by CMs operating in Multiple Transmit Channel Mode.

6.2.4.5 Queue-depth Based Request Frame

The Queue-depth Based Request Frame is the mechanism that a cable modem uses to request bandwidth in terms of bytes, not including or assuming any physical layer overhead (preamble, FEC, physical layer padding, guard time), which is used when the CM is in Multiple Transmit Channel Mode. This is unlike the Request Frame in which requests are made in units of mini-slots that include physical layer overhead. The Queue-depth Based Request Frame is only applicable in the upstream. The CM MUST NOT include any Data PDUs following the Queue-depth Based Request Frame. The CM MUST comply with figure 6-9 and table 6-9 for Queue-depth Based Request Frames.

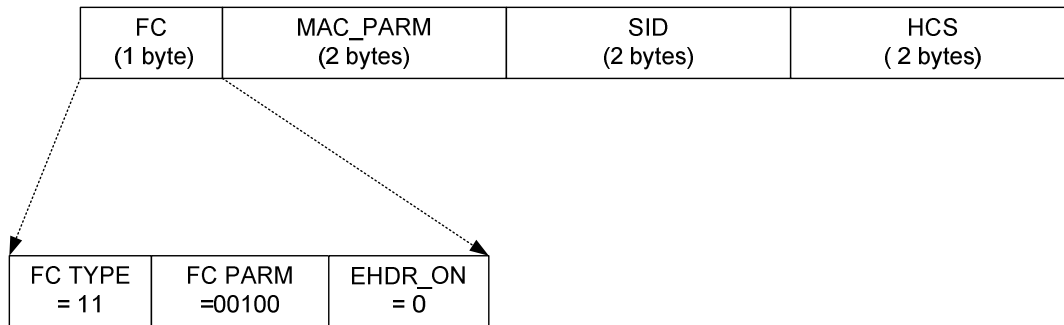


Figure 6-9: Queue-depth Based Request Frame Format

Table 6-9: Queue-depth Based Request Frame Format

Field	Usage	Size
FC	FC_TYPE = 11; MAC-Specific Header FC_PARM[4:0] = 00100; MAC Header only; no data PDU following EHDR_ON = 0; No EHDR allowed	1 byte
MAC_PARM	Total number of bytes requested in units of N bytes, where N is a parameter of the service flow for which this request is being made	2 bytes
SID	Service ID (0...0x3DFF)	2 bytes
EHDR	Extended MAC Header not allowed	0 byte
HCS	MAC Header Check Sequence	2 bytes
	Length of a Queue-depth Based REQ MAC Header	7 bytes

Because the Queue-depth Based Request Frame does not have a Data PDU following it, the LEN field is not needed. The CM MUST replace the LEN field with a SID. The SID uniquely identifies a particular Service Flow within a given CM.

NOTE: The Queue-depth Based Request Frame is one byte longer than the Pre-3.0 DOCSIS Request Frame.

Queue-depth Based Request Frames MUST NOT be used by CMs operating with Multiple Transmit Channel Mode disabled.

6.2.4.6 Concatenation Header

A Specific MAC Header is defined to allow multiple MAC frames to be concatenated.

The CM MUST comply with figure 6-10 and table 6-10 for Concatenation MAC Headers.

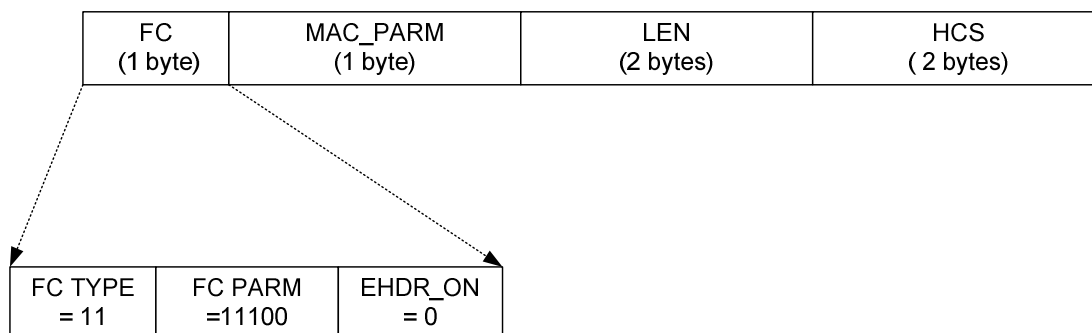


Figure 6-10: Concatenation MAC Header Format

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 byte
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 indicates the total count of MAC Frames (CNT) in this concatenation burst.

The Concatenation Frame is for Pre-3.0 DOCSIS support and MUST NOT be used by CMs operating in Multiple Transmit Channel Mode.

6.2.5 Extended MAC Headers

Every MAC Header, except the Timing, Concatenation MAC Header, Request Frame and Queue-depth Based Request Frame, has the capability of defining an Extended Header field (EHDR). The CM or CMTS MUST indicate the presence of an EHDR field by the EHDR_ON flag in the FC field being set. Whenever this bit is set, then the CM or CMTS MUST use the MAC_PARM field as the EHDR length (ELEN). The minimum defined EHDR is 1 byte. The maximum EHDR length is 240 bytes.

A compliant CMTS and CM MUST support extended headers.

The CM MUST comply with figure 6-11 and table 6-11 for MAC Headers with an Extended Header. The CMTS MUST comply with figure 6-11 and table 6-11 for MAC Headers with an Extended Header.

NOTE: The CM MUST NOT use Extended Headers in a Concatenation MAC Header, but may be included as part of the MAC Headers within the concatenation.

The CM MUST NOT use Extended Headers in Request Frames or Queue-depth Based Request Frames. The CM and CMTS MUST NOT use Extended Headers in Timing MAC Headers.

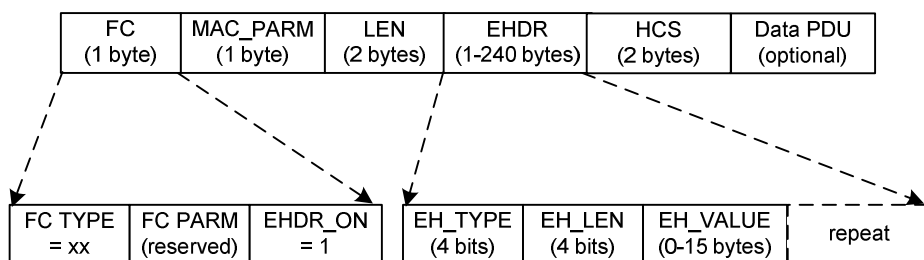


Figure 6-11: Extended MAC Format

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 in 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 CM or CMTS MUST increase the value of the LEN field 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. The size of each EH element is variable. The CM or CMTS MUST set the first byte of the EH element to contain a type and a length field. Every CM MUST use this length to skip over any unknown EH elements. The CM MUST comply with table 6-12 for EH elements. The CMTS MUST comply with table 6-12 for EH elements.

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 byte to 15 bytes

The CM MUST support the types of EH element defined in table 6-13. The CMTS MUST support the types of EH element defined in table 6-13. The CM MUST comply with table 6-13 for Extended Header Types. The CMTS MUST comply with table 6-13 for Extended Header Types. Reserved and extended types are undefined at this point and MUST be ignored by CMs and CMTSs.

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 left attached by the CMTS 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 used by CMs and CMTSs 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 is 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 [12]
	5	Upstream Privacy with Fragmentation EH Element (see [12] and clause 7.2.5.2)
4 (= BP_DOWN)	4	Downstream Privacy EH Element [12]
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 (= BP_UP2)	3	Upstream Privacy EH version 2 Element with no piggyback request
8	varies	Downstream Service EH Element
9	5	DOCSIS Path Verify EH Element
10 to 14		Reserved [CM ↔ CM]
15	XX	Extended EH Element: EHX_TYPE (1 byte), EHX_LEN (1 byte), EH_VALUE (length determined by EHX_LEN)
NOTE: An Upstream Privacy with Fragmentation EH Element only occurs within a Fragmentation MAC-Specific Header (refer to clause 6.2.5.4).		

6.2.5.1 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 clause 7.2.2).

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

6.2.5.2 Request Extended Header

The Request Extended Header (EH_TYPE=1) is used to piggyback requests on packets that do not have the Baseline Privacy extended headers. In that case, when operating with Multiple Transmit Channel Mode disabled, the CM MUST use either the Request Extended Header with EH_LEN=3 or the BP_UP Extended Header to send piggyback requests. When the CM is operating with Multiple Transmit Channel Mode enabled and segment headers are disabled, the CM MUST NOT use piggyback requests. When the CM is operating with Multiple Transmit Channel Mode enabled and segment headers are enabled, the CM MUST only use the request field in the segment header to send a piggyback request.

6.2.5.3 Fragmentation Extended Header

Pre-3.0 DOCSIS fragmented packets use a combination of the Fragmentation MAC header and a modified version of the Upstream Privacy Extended header. Clause 6.2.5.4 describes the Fragmentation MAC header. The Upstream Privacy Extended Header with Fragmentation, also known as the Fragmentation Extended Header, transmitted by the CM MUST comply with table 6-14. CMs operating in Multiple Transmit Channel Mode MUST NOT use fragmentation extended headers.

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 [12]	1 bit
	SID; Service ID associated with this fragment	14 bits
	REQ; number of mini-slots for a piggyback request	8 bits
	Reserved; 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

6.2.5.4 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.

6.2.5.4.1 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 both PHSI-indexed Payload Header Suppression and DSID-indexed Payload Header Suppression. A CMTS MUST support PHSI-indexed Payload Header Suppression. A CMTS SHOULD support DSID-indexed Payload Header Suppression.

The CM MUST comply with table 6-15 for Payload Header Suppression Extended Header sub-elements. The CMTS MUST comply with table 6-15 for Payload Header Suppression Extended Header sub-elements.

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	8 bits
	1 to 254	
	255	

For PHSI-indexed PHS the Payload Header Suppression Index is unique per service flow 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).

For DSID-indexed PHS, the EH_VALUE field of the Payload Header Suppression EHDR is set to the static value of 255. Payload Header Suppression is disabled if this Extended Header element is omitted or, if included, with the EH_Value field set to 0. In DSID-indexed PHS, the DSID references the Payload Header Suppression Field (PHSF).

The CM MUST begin the Upstream Suppression Field with the first byte following the MAC Header Checksum. The CMTS MUST begin the Downstream Suppression Field 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 [12]) 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 by the CM or CMTS in the CRC calculation.

6.2.5.4.2 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 clause 7.2.3.3).

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 by the CM in the Extended Header Element. 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
	1 to 254	Payload Header Suppression Index (PHSI)	(always present)
	Queue Indicator		1 bit
	Active Grants		7 bits

6.2.5.5 BP_UP2 Extended Header

The BP_UP2 EHDR is used when Multiple Transmit Channel Mode is enabled and Baseline Privacy is enabled. When Multiple Transmit Channel Mode is enabled for the CM and segment headers are enabled for a given service flow, the CM MUST use the piggyback opportunity in the segment header for any piggyback requests for that service flow. If segment headers are not enabled for a service flow, the CM is not permitted to create piggyback requests for that service flow. Thus, a piggyback field is not needed in the BP_UP2 EHDR for any service flows. The CM operating in Multiple Transmit Channel Mode with Baseline Privacy Enabled MUST use the BP_UP2 EHDR with a length of 3 for all service flows. The CM MUST comply with table 6-17 for the BP_UP2 EHDR with length of 3.

Table 6-17: BP_UP2 EHDR with Length 3

EH Element Fields	Usage	Size
EH_TYPE	Upstream Privacy EH_TYPE = 7	4 bits
EH_LEN	Length of EH_VALUE = 3	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 [12]	1 bit
	Reserved, set to zero	14 bits

6.2.5.6 Downstream Service Extended Header

The Downstream Service Extended Header (DS EHDR) communicates to the CM information on how to process downstream packets. The DS EHDR contents vary depending on the EH_LEN, which may be one, three or five bytes. The CMTS MUST comply with tables 6-18, 6-19 and 6-20 for DS EHDRs. This header is ignored by CMs which do not implement Downstream Channel Bonding.

Table 6-18: One-byte DS EHDR Sub-Element Format

EH Element Fields	Usage	Size
EH_TYPE	Downstream Service EH_TYPE = 8	4 bits
EH_LEN	1	4 bits
EH_VALUE	Traffic Priority	3 bits
	Reserved	5 bits

Table 6-19: Three-byte DS EHDR Sub-Element Format

EH Element Fields	Usage	Size
EH_TYPE	Downstream Service EH_TYPE = 8	4 bits
EH_LEN	3	4 bits
EH_VALUE	Traffic Priority	3 bits
	Reserved	1 bit
	Downstream Service ID (DSID)	20 bits

Table 6-20: Five-byte DS-EHDR Sub-Element Format

EH Element Fields	Usage	Size
EH_TYPE	Downstream Service EH_TYPE = 8	4 bits
EH_LEN	5	4 bits
EH_VALUE	Traffic Priority	3 bits
	Sequence Change Count	1 bit
	Downstream Service ID (DSID)	20 bits
	Packet Sequence Number	16 bits

When the CMTS classifies a packet to a service flow with a nonzero Traffic Priority (see clause C.2.2.5.1), it MUST add a DS EHDR and set the Traffic Priority sub-element to the value of the service flow's Traffic Priority parameter.

When the CMTS transmits a packet from a Group Service Flow assigned to a single downstream channel (i.e. non-bonded), it MUST include a three-byte DS EHDR with a DSID. Refer to clause 9.2.2.

When the CMTS transmits a packet from a Service Flow assigned to a Downstream Bonding Group, the CMTS MUST include a five-byte DS EHDR (except if there is a vendor specific configuration to permit the Service Flow to send non-sequenced packets). The DSID in a five-byte DS EHDR is a Resequencing DSID, which identifies a resequencing context. The Packet Sequence Number identifies the sequence number of a packet within the resequencing context identified by the DSID.

A Sequenced Null Packet is defined as a variable-length packet-based MAC frame (clause 6.2.2.1) which includes a five-byte Downstream Service EHDR, does not include any other Extended Header and has a Packet PDU length of zero. A CMTS MAY send Sequenced Null Packets. A CM MUST accept Sequenced Null Packets.

For a Resequencing DSID, a packet received with a 3-byte DS EHDR MUST be processed by the CM as a non-sequenced packet. For a non-resequencing DSID, a packet received with 5-byte DS EHDR MUST be processed by the CM as a non-sequenced packet. A packet received with a 2-byte DS EHDR MUST be treated by the CM identically to the 1-byte DS EHDR (the extra byte is ignored). A packet received with a 4-byte DS EHDR MUST be treated by the CM identically to the 3-byte DS EHDR (the extra byte is ignored). A packet received with a 6-byte or greater DS EHDR MUST be treated by the CM identically to the 5-byte DS EHDR (the extra byte(s) are ignored).

6.2.5.7 DPV Extended Header

Table 6-21: DPV Extended Header Format

EH Element Fields	Usage	Size
EH_TYPE	DPV EHDR = 9	4 bits
EH_LEN	Length of EH_VALUE = 5 bytes	4 bits
EH_VALUE	Start Reference Point	8 bits
	Timestamp Start	32 bits
Start Reference Point:	This is the DPV Reference Point that the DPV measurement originates from (see clause 10.4.2).	
Timestamp Start:	This is the local timestamp at the sender when the DPV packet gets injected into the data stream and departs from the DPV reference point.	

The CMTS MAY support the generation of the DPV Extended Header. The CMTS MAY place a DPV EHDR on any packet within any DSID or any Service Flow. The CMTS MUST comply with table 6-21 for DPV EHDRs. A Modular CMTS Core MAY choose to place a DPV EHDR on any packet within any DEPI flow. This may be done in order to compare the average latency between different Service Flows and/or DEPI flows.

The CM MAY support the generation of the DPV Extended Header.

The CMTS and CM are not required to take any action upon receiving a DPV EHDR other than silently discarding it.

6.3 Segment Header Format

The CM MUST use a Segment Header when transmitting packets in Multiple Transmit Channel Mode for service flows where use of the segment header is enabled. For these service flows, a Segment Header must appear at the beginning of any transmission made with IUCs 5,6,9,10 or 11. Figure 6-12 shows the segment header format. The segment header is 8 bytes in length. table 6-22 describes the segment header fields. The CM MUST comply with figure 6-12 and table 6-22 for segment headers.

PFI (1 bit)	R (1 bit)	Pointer Field (14 bits)	Sequence # (13 bits)	SC (3 bits)	Request (2 Bytes)	HCS (2 Bytes)
----------------	--------------	----------------------------	-------------------------	----------------	----------------------	------------------

Figure 6-12: Segment Header Format

Table 6-22: Segment Header Fields

Field	Usage	Size
PFI	Pointer Field Indicator. This bit is set to a one, to indicate that the pointer field is relevant. When cleared to a zero, this bit indicates that there is no DOCSIS MAC frame starting within this segment and the pointer field is ignored.	1 bit
R	Reserved. This field should be set to a zero by the CM.	1 bit
Pointer Field	When the PFI bit is a one, the value in this field is the number of bytes past the end of the segment header that the receiver will skip when looking for a DOCSIS MAC Header. Thus, a value of zero in the pointer field with the PFI set to one would designate a DOCSIS MAC header beginning just after the segment header.	14 bits
Sequence #	Sequence number that increments by 1 for every segment of a particular service flow.	13 bits
SC	SID Cluster ID of the SID Cluster associated with the Request field of the segment header. The valid SID Cluster ID range is 0 to M-1, where M is the number of SID Clusters per Service Flow supported by the CM.	3 bits
Request	The total number of bytes requested in units of N bytes where N is a parameter of the service flow for which the request is being made. See clause C.2.2.6.12.	2 bytes
HCS	MAC Header Check Sequence. Similar to HCS used on all MAC headers and is calculated over all other fields in the segment header.	2 bytes

The HCS field is a 16-bit CRC that ensures the integrity of the segment header, even in a collision environment. The CM MUST include all fields within the segment header for the HCS field coverage except the HCS field itself. The HCS is calculated using CRC-CCITT ($x^{16} + x^{12} + x^5 + 1$) as defined in [21].

For segment header ON operation, the CM may use the piggyback field in the segment header to make piggyback requests for the service flow and MUST NOT use any request EHDR fields within the segment payload.

6.4 MAC Management Messages

6.4.1 MAC Management Message Header

CMs and CMTSs **MUST** encapsulate MAC Management Messages in an LLC unnumbered information frame per [17], which in turn is encapsulated within the cable network MAC framing, as shown in figure 6-13. Figure 6-13 shows the MAC Header and the MAC Management Message Header fields which are common across all MAC Management Messages.

The CMTS **MUST** use a unique MAC address for each MAC Domain interface. This address is used by the CMTS as the Source Address for all MAC Management Messages for the MAC Domain. Since the CM is required to use the Source Address of the MDD messages to identify channels associated with the MAC Domain of its Primary DS channel, topology resolution (clause 10.2.3.3) could fail if multiple MAC Domains use the same MAC address and have DS channels which reach the same CM.

The CMTS **MUST NOT** add a Downstream Service EHDR to the following MAC Management Message types: SYNC, UCD (types 2, 29 or 35), MAP, DCD and MDD. The CMTS **MAY** add a three-byte Downstream Service EHDR to any other type of MAC Management Message. If this EHDR is present, the CM **MUST** filter the MAC Management Message in accordance with the rules of clause 9.2.2.4. The CM **MUST NOT** forward MAC Management Messages to any interface or eSAFE.

DOCSIS 3.0 does not define support for sequenced downstream MAC Management Messages. A CMTS **MUST NOT** transmit a MAC Management Message with a five-byte Downstream Service Extended Header. A CM **MUST** silently discard a MAC Management Message containing a five-byte Downstream Service Extended Header. This does not preclude future versions of this specification from defining sequenced MAC Management Messages using a five-byte Downstream Service Extended Header.

The CMTS **MUST NOT** add a Service Flow EHDR to MAC Management Messages. The CM **MUST NOT** add a Service Flow EHDR to MAC Management Messages.

See [15] for rules governing the use of the Baseline Privacy EHDR on MAC Management Messages.

Unless otherwise specified, a CMTS may transmit and a CM **MUST** accept a downstream MAC Management Message to the CM's individual MAC address on any downstream channel received by the CM.

Unless otherwise specified, a CM may send and a CMTS **MUST** accept, an upstream MAC Management Message on any upstream channel transmitted by the CM.

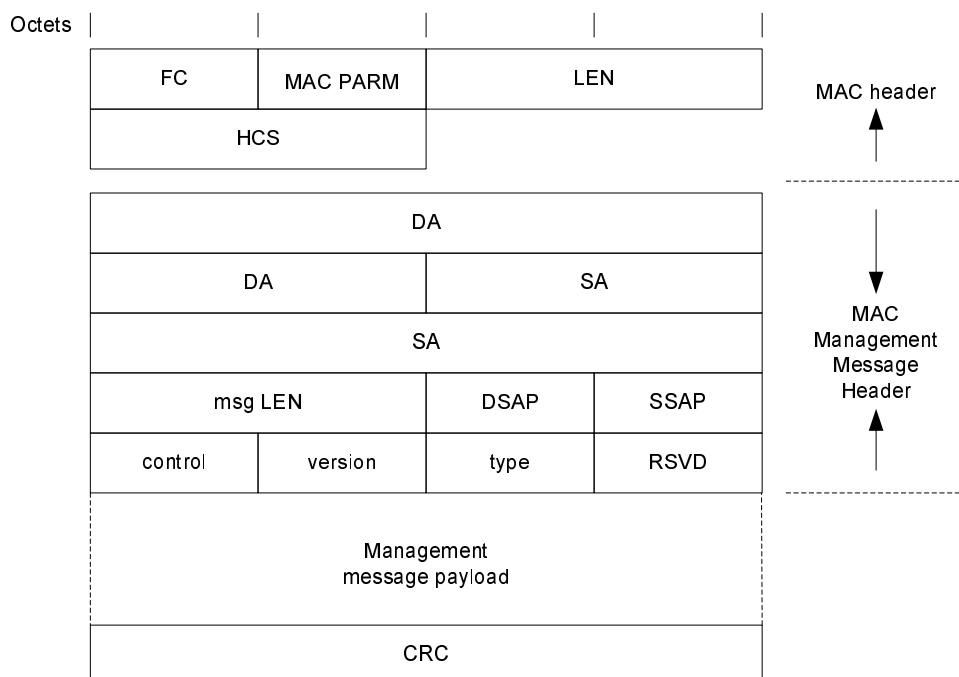


Figure 6-13: MAC Header and MAC Management Message Header Fields

The fields of the MAC Management Message Header shown in figure 6-13 are defined below:

FC, MAC PARM, LEN, HCS, Common MAC frame header: Refer to clause 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 annex A.

Source Address (SA): The MAC address of the source CMTS MAC Domain Interface or source CM.

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 [17].

SSAP: The LLC null source SAP (00) as defined by [17]. Set to 0 for this version for all messages other than the RNG-REQ, INIT-RNG-REQ and B-INIT-RNG-REQ. See clause 6.4.5.

Control: Unnumbered information frame (03) as defined by [17].

Version and Type: Each 1 octet. Refer to table 6-23. Messages with a version number of 1 are understood by all CMs and CMTSs compliant with all versions of the DOCSIS specification. Messages with a version number of 2 are understood by DOCSIS 1.1, 2.0 and 3.0 equipment. Messages with a version number of 3 are understood by DOCSIS 2.0 and 3.0 equipment. Messages with a version number of 4 are understood by DOCSIS 3.0 equipment. DOCSIS 3.0 compliant CMs and CMTSs MUST silently discard any message with version number greater than 4.

Table 6-23: MAC Management Message Types

Type Value	Version	Message Name	Message Description
1	1	SYNC	Timing Synchronization
2, 29 or 35	1, 3 or 4	UCD	Upstream Channel Descriptor A UCD for a DOCSIS 3.0 Only channel uses a type of 35 and a version of 4. A UCD for a DOCSIS 2.0/3.0 Only Channel uses a type of 29 and a version of 3. All other UCDs use a type of 2 and a version of 1 (see clause 6.4.3).
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		Reserved (deprecated)
11	1		Reserved (deprecated)
12	1	BPKM-REQ	Privacy Key Management Request [15]
13	1	BPKM-RSP	Privacy Key Management Response [15]
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	3		(See entry for UCD above)
30	3	INIT-RNG-REQ	Initial Ranging Request
31	3	TST-REQ	Test Request Message
32	3	DCD	Downstream Channel Descriptor [5]
33	4	MDD	MAC Domain Descriptor
34	4	B-INIT-RNG-REQ	Bonded Initial Ranging Request
35	4		(See entry for UCD above)
36	4	DBC-REQ	Dynamic Bonding Change Request
37	4	DBC-RSP	Dynamic Bonding Change Response
38	4	DBC-ACK	Dynamic Bonding Change Acknowledge
39	4	DPV-REQ	DOCSIS Path Verify Request
40	4	DPV-RSP	DOCSIS Path Verify Response
41	4	CM-STATUS	Status Report
42	4	CM-CTRL-REQ	CM Control
43	4	CM-CTRL-RSP	CM Control Response
44	4	REG-REQ-MP	Multipart Registration Request
45	4	REG-RSP-MP	Multipart Registration Response
46 to 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 of DOCSIS for all messages other than the RNG-REQ and INIT-RNG-REQ. See clause 6.4.5.

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

CRC: Covers message including header fields (DA, SA, etc.). Polynomial defined by [17].

A compliant CMTS or CM MUST support the MAC management message types listed in table 6-23.

6.4.2 Time Synchronization (SYNC)

Time Synchronization (SYNC) MUST be transmitted by CMTS at a periodic interval to establish MAC sublayer timing. The CMTS MUST format this message to use an FC field with FC_TYPE = MAC Specific Header and FC_PARM = Timing MAC Header, followed by a Packet PDU in the format shown in figure 6-14.

The CMTS MUST transmit SYNCs on Primary-Capable DS Channels. The CMTS MUST NOT transmit SYNCs on non-Primary Capable DS Channels.

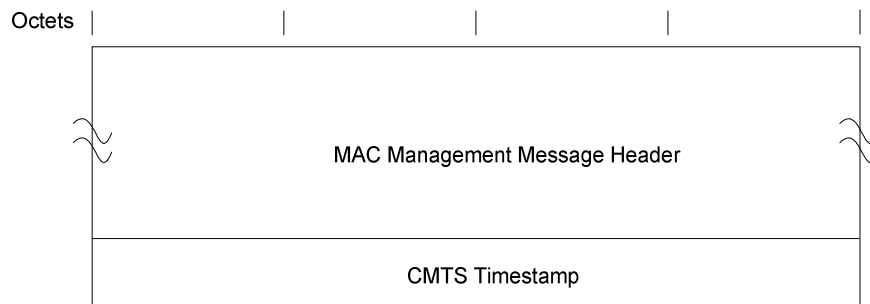


Figure 6-14: Format of Packet PDU Following the Timing Header

The parameters shall be as defined below:

CMTS Timestamp: The count state of an incrementing 32-bit binary counter clocked with the CMTS 10,24 MHz master clock.

The CMTS timestamp represents the count state at the instant that the first byte (or a fixed time offset from the first byte) of the Time Synchronization MAC Management Message is transferred from the Downstream Transmission Convergence Sublayer to the Downstream Physical Media Dependent Sublayer as described in [4]. The CMTS MUST NOT allow a SYNC message to cross an MPEG packet boundary.

6.4.3 Upstream Channel Descriptor (UCD)

An Upstream Channel Descriptor MUST be transmitted by the CMTS at a periodic interval to define the characteristics of a logical upstream channel (figure 6-15). A separate message MUST be transmitted by the CMTS for each logical upstream that is currently available for use. The CMTS MUST send UCD messages for a given upstream channel on the same downstream channel(s) that it sends the MAP messages for that upstream channel.

The following table describes the linkage between channel types, UCD types, logical channel types, burst descriptor types and the DOCSIS modes in which a CM is able to use the channel. The table also indicates the clauses of the specification in which the particular item is detailed.

Table 6-24: Linkage Between Channel Types

Upstream Channel Type	Channel Description	UCD Type/Version (clause 6.4.3)	Logical Channel Types (clause 6.1.2.6)	Burst Descriptors (clause 6.4.3)	Usable by CMs in DOCSIS Operational Mode
Type 1	DOCSIS 1.x PHY Channel	2/1	Type 1	DOCSIS 1.x only (Type 4)	DOCSIS 1.x, 2.0, 3.0
Type 2	Mixed DOCSIS 1.x/2.0 TDMA PHY channel	2/1	Type 2	DOCSIS 1.x and 2.0 (Type 4 for 1.x and Type 5 for 2.0 TDMA)	DOCSIS 1.x, 2.0, 3.0
Type 3	DOCSIS 2.0 PHY channel	29/3	Type 3A (2.0 TDMA) or Type 3S (2.0 S-CDMA)	DOCSIS 2.0 (Type 5)	DOCSIS 2.0 and 3.0
Type 4	DOCSIS 3.0 PHY channel	35/4 and 29/3	Type 4A (2.0 or 3.0 TDMA) or Type 4S (2.0 or 3.0 S-CDMA)	DOCSIS 2.0 (Type 5)	DOCSIS 3.0 and 2.0
		35/4	Type 4AR (3.0 TDMA) or Type 4SR (3.0 S-CDMA)	DOCSIS 2.0 (Type 5)	DOCSIS 3.0

The MAC management header for this message has 3 possible values for the Type field and for the Version field. For a Type 4 channel, the CMTS MUST use a value of 35 for the Type field and use a value of 4 for the Version field. For Type 3 channels, the CMTS MUST use a value of 29 for the Type field and a value of 3 for the Version field. For Type 1 and Type 2 channels, the CMTS MUST use a value of 2 for the Type field and a value of 1 for the Version field.

Depending on the IUC UCD message Type and Channel Type, burst descriptors can be encoded as either Type 4 or Type 5 TLVs. A CMTS MUST NOT use Type 5 TLVs to encode IUCs 1-6 in a UCD with a message Type of 2. If a Type 2 UCD describes a mixed 1.x/2.0 PHY logical channel, the CMTS MUST additionally contain Type 5 TLV burst descriptors for IUCs 9 and/or 10 and/or 11 providing advanced TDMA data opportunities in the UCD. Advanced TDMA burst descriptor attributes are those that can be included in a Type 5 burst descriptor but cannot be included in a Type 4 burst descriptor. A CMTS MUST use only Type 5 TLVs to encode burst profiles in a UCD with a message Type of 29. A CMTS MUST use only Type 5 TLVs to encode burst profiles in a UCD with a message Type of 35.

A Type 29 UCD transmitted by a CMTS MUST contain a Type 5 burst descriptor for ranging, a Type 5 burst descriptor for requests and a Type 5 burst descriptor for data.

In a Type 29 UCD a CMTS MUST NOT include burst descriptors for IUCs 5 or 6 in a UCD message for a Type 3 Upstream Channel.

In a Type 35 UCD a CMTS MUST include burst descriptors for data grants corresponding to IUCs 5, 6, 9 and 10.

For a Type 35 UCD, the CMTS MAY include:

- Burst attributes that enable SAC Mode 2 and Code Hopping Mode 2.
- Burst attributes associated with IUC 11 that are not intended for UGS.

To make use of the UCD possibilities enumerated above, the channel described by the Type 35 UCD can only be used by DOCSIS 3.0 CMs.

A channel could be shared by DOCSIS 3.0 and DOCSIS 2.0 CMs using a UCD of Type 29 and a UCD of Type 35 to describe the same channel corresponding to the same Upstream Channel ID. However, only one set of MAPs pertaining to the UCID is generated and grants are allocated in the MAP. The purpose of this multiple UCD to UCID mapping is for conservation of a logical channel in the case DOCSIS 2.0 CMs and DOCSIS 3.0 CMs operate in the same frequency channel. Because a UCD of Type 29 is not allowed to have burst descriptors for IUC 5 and 6, using a Type 29 UCD for both DOCSIS 2.0 and 3.0 CMs restricts the DOCSIS 3.0 CMs operating in Multiple Transmit Channel Mode from being commanded to use burst profiles for data transmissions from up to five assigned burst profiles in the UCD message (IUC 5, 6, 9, 10 and 11). Assigning the DOCSIS 2.0 CMs and 3.0 CMs to separate logical channels is a solution subject to disadvantages of loss of statistical multiplexing gain and consumption of a logical channel resource at the CMTS.

For a channel that is described using a UCD of Type 29 and a UCD of Type 35 the CMTS MUST send UCDs that comply with the following:

- Transmission parameters like Minislot size, Modulation rate, Preamble pattern etc. are identical in each of the UCDs in the set.
- Burst attributes corresponding to the same IUC are identical in each of the UCDs in the set (if the corresponding burst profile is present in both UCDs).
- The Configuration Change Count of each UCD is identical and matches the UCD Configuration Change Count in the MAP.
- The UCD29 includes a Type 22 TLV with a value of 1.
- The UCD35 includes a Type 20 TLV with a value of 0 or 1.

When a CM is commanded to another upstream channel without specific UCD configuration information (e.g. in the case of upstream channel override or in the case of a DCC or DBC request without UCD configuration information), the CM MUST look for UCDs containing the assigned UCID in the active downstreams and select from the existing UCDs the UCD with the highest Type value consistent with the CM's capability. In other words, the CM does not necessarily use the first UCD corresponding to the assigned UCID that it sees. After receiving UCD messages, the CM MUST use the TLV22 bitmap in the UCD (if present) to check if there is another UCD for this UCID with a higher Type value consistent with the CM's capability. In the case when UCD configuration information is provided in the DCC or DBC Request, the CM uses the UCD configuration information immediately. Similarly, if a CM is acquiring a UCD in preparation for ranging on a saved upstream channel, after a reinitialize MAC event, the CM MUST obtain the UCD containing the saved UCID with the highest Type value consistent with the CM's capability.

For interoperability, a CMTS SHOULD provide:

- Burst descriptors for IUCs 1, 5 and 6 in a Type 2 UCD describing a Type 1 channel.
- Burst descriptors for IUCs 1, 5, 6, 9 and 10 in a Type 2 UCD describing a Type 2 channel.
- Burst descriptors for IUCs 1, 9 and 10 in a Type 29 UCD.

Type 4 burst descriptors indicate that the preamble of the burst is in accordance to DOCSIS 1.x specifications while Type 5 burst descriptors indicate that the preamble of the burst is in accordance to DOCSIS 2.0 preambles. In particular, preambles for bursts described by Type 4 burst descriptors are sent using the same modulation as that described for the burst. Preambles for bursts described by Type 5 burst descriptors are sent using either QPSK0 or QPSK1 constellations.

A CMTS MUST consider an upstream as a Type 4 Upstream if:

- the Selectable Active Codes Mode 2 and Code Hopping Mode 2 features are enabled;
- IUCs 5 and 6, are associated with Type 5 burst descriptors; or
- burst attributes associated with IUC 11 are not intended for UGS (though a CMTS can provide UGS opportunities using IUC 11 on a Type 4 Upstream).

The CMTS MUST NOT consider an upstream as a Type 1 or Type 2 Upstream if any of the following is true about the channel wide parameters:

- S-CDMA mode is enabled;
- the Mini-slot size is 1 time tick; or
- the value of the Modulation Rate parameter is 32.

The CMTS MUST NOT consider an upstream as a Type 1 or Type 2 Upstream if any of the following is true about any of IUCs 1-4:

- a modulation type other than QPSK or 16-QAM is used;
- the FEC Error Correction (T) parameter is greater than 10;
- any portion of the extended preamble is used; or
- any attribute from table 6-26 with a Type greater than 11 is present in the descriptor.

A CM MUST be able to recognize Channel Parameter TLVs with Type 20 and 21 even if the CM is not capable of Selectable Active Codes mode 2 and Code Hopping mode 2. If a CM does not support Selectable Active Codes mode 2 and Code Hopping mode 2, then the CM MUST NOT use a UCD that indicates that these features are active.

To provide for flexibility, the message parameters following the Downstream channel ID MUST be encoded by the CMTS in a type/length/value (TLV) form in which the type and length fields are each 1 octet long.

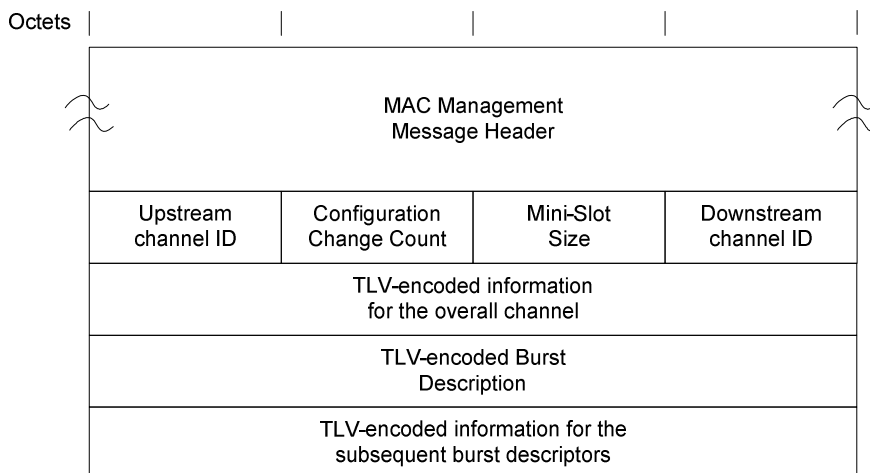


Figure 6-15: Upstream Channel Descriptor

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

Configuration Change Count: Incremented by one (modulo the field size) by the CMTS whenever any of the values of this channel descriptor change, excluding the S-CDMA snapshot TLV. If the value of this count in a subsequent UCD remains the same, the CM can quickly decide that the channel operating parameters 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 microseconds. For channels that can support DOCSIS 1.x CMs, the allowable values are $T = 2^M$, $M = 1, \dots, 7$. That is, $T = 2, 4, 8, 16, 32, 64$ or 128 . For DOCSIS 2.0 or 3.0 Only Channels the relationship between M and T remains the same, but the allowable values are $M = 0, 1, 7$, with $T = 1, 2, 4, 8, 16, 32, 64$ or 128 . If the value of T is 1 then the channel will be treated as a DOCSIS 2.0/3.0 Only Channel. On S-CDMA channels, this parameter will not have any effect.

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

NOTE 1: Upstream Channel ID = 0 is reserved for network management purposes [10].

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

NOTE 2: Downstream Channel ID = 0 is reserved for network management purposes [10].

All other parameters are coded as TLV tuples. The Type values used by the CMTS MUST be those defined in table 6-25, for channel parameters and table 6-26, for upstream physical layer burst attributes. The CMTS MUST place burst descriptors (Type 4 and/or Type 5) that appear in the UCD message after all other channel-wide parameters.

Table 6-25: Channel TLV Parameters

Name	Type (1 byte)	Length (1 byte)	Value (Variable length)
Modulation Rate	1	1	Multiples of base rate of 160 kHz. For TDMA channels, valid Values are 1, 2, 4, 8, 16 or 32. A value of 32 means that this is a DOCSIS 2.0/3.0 Only Upstream. If S-CDMA mode is enabled then the only valid Values for this parameter are 8, 16 and 32.
Frequency	2	4	Upstream center frequency (Hz).
Preamble Pattern	3	1-128	The Value field defines the first portion of the Preamble Superstring. If there is no Extended Preamble Pattern parameter, then this parameter defines the entire Preamble Superstring. All burst-specific preamble values are chosen as bit-substrings of the Preamble Superstring. 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 (DOCSIS 1.x)	4	n	May appear more than once; described below.
Burst Descriptor (DOCSIS 2.0/3.0)	5	n	May appear more than once; described below.
Extended Preamble Pattern	6	1-64	512 Bit Preamble Superstring extension. The Value field is concatenated to the end of the Value field of the Preamble Pattern to complete the Preamble Superstring. This Parameter will not be included unless the length of the Preamble Pattern parameter is 128 bytes. Therefore the MSB of the first byte of the Value field of this parameter always follows the LSB of the 128th byte of the Value field of the Preamble Pattern parameter in the Preamble Superstring.
S-CDMA Mode Enable (see note 1)	7	1	1 = on; 2 = off. If parameter is on, the upstream will operate in S-CDMA mode. Otherwise it operates in TDMA mode. If this parameter is set to on, this is a DOCSIS 2.0/3.0 Only Upstream.
S-CDMA Spreading Intervals per frame	8	1	Number of consecutive spreading intervals mapped onto a two-dimensional frame. (Value is 1 through 32). The CMTS MUST include this TLV only if S-CDMA Mode is enabled for the channel.
S-CDMA Codes per Mini-slot	9	1	Number of consecutive codes mapped into a two-dimensional mini-slot. (Value is 2 through 32). The CMTS MUST include this TLV if and only if S-CDMA Mode is enabled for the channel.
S-CDMA Number of Active Codes	10	1	Number of codes available to carry data payload. (Value is 64 through 128). This value is a multiple of Codes per Mini-slot (TLV type 9). The CMTS MUST include this TLV if and only if S-CDMA Mode is enabled for the channel.
S-CDMA Code Hopping Seed	11	2	15-bit seed to initialize code hopping sequence. The value is left-justified in the 2-byte field. Set seed = 0 to disable code hopping. The CMTS MUST include this TLV if and only if S-CDMA Mode is enabled for the channel.
S-CDMA US ratio numerator 'M'	12	2	The numerator (M) of the M/N ratio relating the downstream symbol clock to the upstream modulation clock. The CMTS MUST include this TLV if and only if S-CDMA Mode is enabled for the channel. The value of M specified in [4] is used.
S-CDMA US ratio denominator 'N'	13	2	The denominator (N) of the M/N ratio relating the downstream symbol clock to the upstream modulation clock. The CMTS MUST include this TLV if and only if S-CDMA Mode is enabled for the channel. The value of N specified in [4] is used.
S-CDMA Timestamp Snapshot (see note 2)	14	9	Snapshot of the timestamp, mini-slot and S-CDMA frame taken at an S-CDMA frame boundary at the CMTS. A new value is sampled and sent with each UCD message. Refer to [12]. The CMTS MUST include this TLV if and only if S-CDMA Mode is enabled for the channel.

Name	Type (1 byte)	Length (1 byte)	Value (Variable length)
Maintain Power Spectral Density	15	1	1=on; 2=off. If this value is on and the modulation rate is different from the previous UCD, the CM MUST change its transmit power level to keep the power spectral density as close as possible to what it was prior to the modulation rate change unless it is operating in Multiple Transmit Channel Mode. If this value is off or this parameter is omitted, then the CM maintains the same power level that it was using prior to the modulation rate change. If the CM is operating in Multiple Transmit Channel Mode, it MUST ignore this TLV and behave as if the TLV was set to off. In any case the effect of this parameter only lasts until the CM receives a power adjustment in a RNG-RSP.
Ranging Required	16	1	0= no ranging required. 1= unicast initial ranging required. 2= broadcast initial ranging required. If this value is non-zero and the UCD change count does not match the UCD currently in effect, the CM MUST perform ranging as specified by this TLV before using any other transmit opportunities with the new UCD parameters. If ranging is required and the CM is already registered, then it MUST maintain its SIDs and not re-register. If this value is 0 or this TLV is omitted, no ranging is required.
S-CDMA Maximum Scheduled Codes enabled	17	1	1 = Maximum Scheduled Codes is enabled. 2 = Maximum Scheduled Codes is disabled. CMs that implement the S-CDMA Maximum Scheduled Codes MUST set the RSVD field in the Ranging Requests as described in clause 6.4.5. The CMTS MUST NOT include this TLV if S-CDMA Mode is disabled for the channel.
Ranging Hold-Off Priority Field	18	4	Bit Field with values representing device classes, as defined in clause C.1.3.1.16 that should temporarily inhibit Initial Ranging. The CMTS may include this TLV in the UCD message. The CM uses this TLV as described in clause 10.2.3.4.
Channel Class ID	19	4	Bit Field with values representing device classes as defined in clause C.1.3.1.16 that are allowed to use the channel. The CMTS may include this TLV in the UCD message. The CM uses this TLV as described in clause 10.2.3.4.
S-CDMA selection mode for active codes and code hopping	20	1	0 = Selectable active codes mode 1 enabled and code hopping disabled. 1 = Selectable active codes mode 1 enabled and code hopping mode 1 enabled. 2 = Selectable active codes mode 2 enabled and code hopping mode 2 enabled. 3 = Selectable active codes mode 2 enabled and code hopping disabled. The set of active codes is selectable via TLV type 21. The CMTS MUST NOT include this TLV in a Type 2 or 29 UCD. The CMTS MUST include this TLV in a Type 35 UCD if S-CDMA Mode is enabled. The CM MUST ignore this TLV in a Type 2 or 29 UCD or in a Type 35 UCD where S-CDMA Mode is disabled.
S-CDMA selection string for active codes	21	16	128-bit string indicating which codes are active. The first element in the string corresponds to code 0 (the all-ones code) and the last element in the string corresponds to code 127. A "1" element in the string indicates an active code and a "0" indicates an unused code. The CMTS sets the number of ones in the string equal to the S-CDMA Number of Active Codes (TLV type 10). The CMTS MUST include this TLV if TLV encoding type 20 is included in the UCD and has value equal to 2 or 3. The CMTS MUST NOT include this TLV in a Type 2 or Type 29 UCD. The CM MUST ignore this TLV in a Type 2 or Type 29 UCD.
Higher UCD for the same UCID present bitmap	22	1	Bit 0: 1 if UCD35 is present for this UCID; 0 if UCD35 is not present Bits 1-7: Reserved for future use, set to 0.
NOTE 1: CM MUST assume S-CDMA mode is off if TLV is not present.			
NOTE 2: Refer to [12], for a description of the Timestamp Snapshot. A change solely in this parameter for a particular UCD does not represent a change in overall channel operating parameters, hence the UCD channel change count will not be implemented.			

Burst Descriptors are composed of an upstream Interval Usage Code, followed by TLV encodings that define the physical-layer characteristics that are to be used during that interval. The upstream interval usage codes are defined in the MAP message clause of this specification (see clause 6.4.4 and table 6-27). The CMTS MUST comply with figure 6-16 for Burst Descriptors.

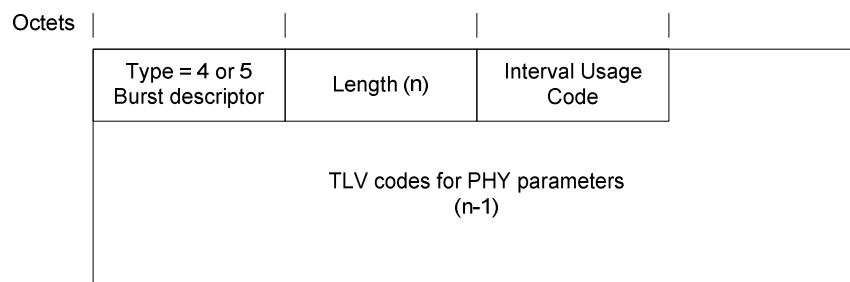


Figure 6-16: Top-Level Encoding for Burst Descriptors

In figure 6-16:

Burst Descriptor: Type 4 Burst Descriptors intended for DOCSIS 1.x and/or DOCSIS 2.0/3.0 modems; Type 5 for Burst Descriptors intended for DOCSIS 2.0/3.0 modems only.

Length: The number of bytes in the overall object, including the IUC and the embedded TLV items.

IUC: Interval Usage code defined in table 6-27. The IUC is coded on the 4 least-significant bits. The 4 most-significant bits are unused (=0).

TLV items: TLV parameters as described in table 6-26.

Two different type values are used to describe Burst Descriptors. Type 4 Burst Descriptors are understood by all modems and are only be used to describe IUCs 1 through 6 from table 6-27. Type 5 Burst Descriptors are understood by DOCSIS 2.0 or 3.0 modems. A type 5 burst descriptor MUST be used by a CMTS to describe any IUC if any of the following is true: a modulation type other than QPSK or 16-QAM is used, the FEC Error Correction (T) attribute is greater than 10, any portion of the Extended Preamble is used or any attribute from table 6-26 with a type greater than 11 is present in the descriptor. Type 5 burst descriptors MUST NOT be used by the CMTS to describe IUC 5 or IUC 6 in a Type 2 UCD.

A Burst Descriptor MUST be included by the CMTS for each Interval Usage Code that is to be used in the allocation MAP. The Interval Usage Code used by the CMTS MUST be one of the values from table 6-27.

Within each Burst Descriptor is an unordered list of Physical-layer attributes, encoded as TLV values. These attributes are shown in table 6-26. The CMTS MUST ensure that the set of burst attributes for all the burst descriptors in the UCD allow any CM not operating in Multiple Transmit Channel Mode on the upstream to be able to request enough mini-slots to be able to transmit a maximum size PDU (see clause 6.2.2).

Table 6-26: Upstream Physical-Layer Burst Attributes

Name	Type (1 byte)	Length (1 byte)	Value (Variable Length)
Modulation Type	1	1	1 = QPSK. 2 = 16-QAM. 3 = 8-QAM. 4 = 32-QAM. 5 = 64-QAM. 6 = 128-QAM (S-CDMA only). Values greater than 2 are not used in a descriptor encoded in a type 4 Burst Descriptor.
Differential Encoding	2	1	1 = on, 2 = off (see [12]).
Preamble Length	3	2	Up to 1 536 bits for a type 5 Burst Descriptor. Up to 1 024 bits for a type 4 Burst Descriptor. If this descriptor is encoded in a type 4 TLV, then the substring of the Preamble Superstring defined by this parameter and the Preamble Value Offset MUST NOT include any bits from the Extended Preamble Pattern. The value MUST be an integer number of symbols (see [12]).
Preamble Value Offset	4	2	Identifies the bits to be used in the preamble. This is specified as a starting offset into the Preamble Super string. 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 Superstring. A value of 100 means that the preamble is to use the 101st and succeeding bits from the Preamble Superstring. This value is a multiple of the symbol size. The first bit of the preamble is the first bit into the symbol mapper and is 11 in the first symbol of the burst (see [12]).
FEC Error Correction (T)	5	1	0 to 16 for descriptors encoded in a type 5 Burst Descriptor. 0 to 10 for descriptors encoded in a type 4 Burst Descriptor. (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. The CMTS MUST include this TLV with a value greater than zero when the interval type is Short Data Grant (IUC 5) or Advanced PHY Short Data Grant (IUC 9), (see clause 7.2.1.2.5). 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	For TDMA channels, the number of modulation intervals 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 Type 4 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 increased by 1. For S-CDMA channels, there is no guard time and hence the CM MUST ignore this value. This TLV will not be present for S-CDMA channels.
Last Codeword Length	10	1	1 = fixed; 2 = shortened.
Scrambler on/off	11	1	1 = on; 2 = off.
R-S Interleaver Depth (Ir)	12	1	Reed-Solomon block interleaving depth. A depth of 0 indicates Dynamic Mode; a depth of 1 indicates RS Interleaving Disabled (see [12]) (0 through floor $\lceil 2 \cdot 048 / (K+2T) \rceil$). This TLV MUST be present for burst descriptors encoded in type 5 Burst Descriptors on DOCSIS 2.0/3.0 TDMA channels. This TLV MUST NOT be present for S-CDMA channels or in descriptors encoded in a type 4 Burst Descriptor.
R-S Interleaver Block Size (Br)	13	2	Reed-Solomon block interleaving size in Dynamic Mode. ($2 \cdot N_r$ through 2 048 where $N_r = K+2T$). This TLV MUST be present in burst descriptors encoded in type 5 Burst Descriptors for DOCSIS 2.0/3.0 TDMA channels. This TLV MUST NOT be present on S-CDMA channels or in descriptors encoded in a type 4 Burst Descriptor.

Name	Type (1 byte)	Length (1 byte)	Value (Variable Length)
Preamble Type	14	1	1 = QPSK0. 2 = QPSK1. (Reference [12]). This TLV MUST NOT be present in descriptors encoded in a type 4 Burst Descriptor.
S-CDMA Spreader on/off	15	1	1 = on; 2 = off. This TLV MUST be present for S-CDMA channels. This TLV MUST NOT be present for non-S-CDMA channels or in descriptors encoded in a type 4 Burst Descriptor.
S-CDMA Codes per Subframe	16	1	Number of codes per sub-frame used in the S-CDMA framer. (1 through 128). This TLV MUST be present for S-CDMA channels. This TLV MUST NOT be present for non-S-CDMA channels or in descriptors encoded in a type 4 Burst Descriptor.
S-CDMA Framer Interleaving Step Size	17	1	Size of interleaving steps used in S-CDMA framer. (1 through 31). This TLV MUST be present for S-CDMA channels. This TLV MUST NOT be present for non-S-CDMA channels or in descriptors encoded in a type 4 Burst Descriptor.
TCM Encoding	18	1	1 = on; 2 = off. This TLV MUST be present for S-CDMA channels. This TLV MUST NOT be present for non-S-CDMA channels or in descriptors encoded in a type 4 Burst Descriptor.

6.4.3.1 Example of UCD Encoded TLV Data

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

Type 1	Length 1	Modulation Rate	
Type 2	Length 4	Frequency	
Type 3	Length 1-128	Preamble Pattern	
Type 6	Length 1-64	Extended Preamble Pattern	
Type 4	Length N	First Burst Descriptor	
Type 4	Length N	Second Burst Descriptor	
Type 5	Length N	Third Burst Descriptor	
Type 5	Length N	Fourth Burst Descriptor	

Figure 6-17: Example of UCD Encoded TLV Data

6.4.4 Upstream Bandwidth Allocation Map (MAP)

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

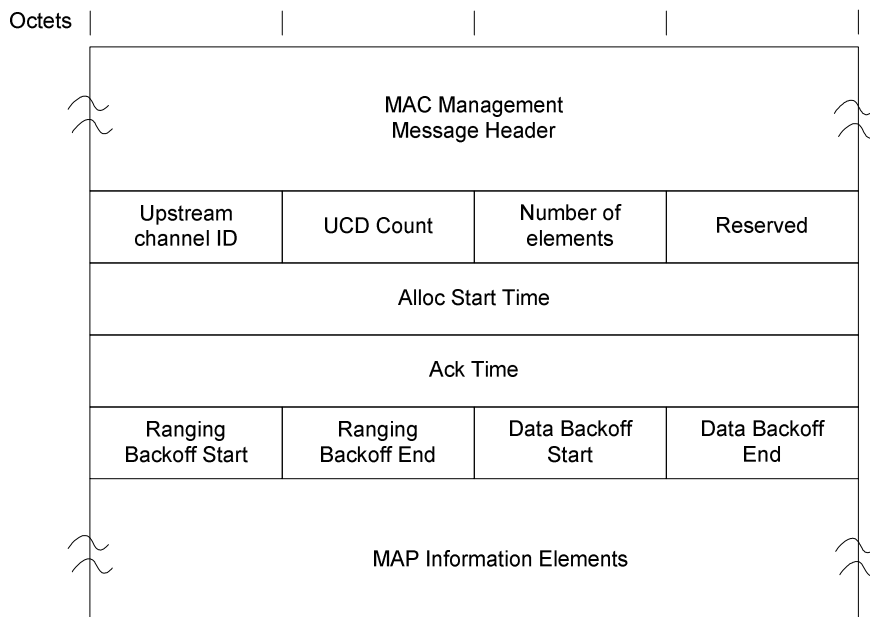


Figure 6-18: MAP Format

The parameters of MAP messages transmitted by a CMTS 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 clause 11.1.

Number of Elements: Number of information elements in the MAP.

Reserved: Reserved field for 32-bit boundary 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 (in mini-slots) processed in the upstream. This time is used by the CMs for collision detection purposes. See clause 7.2.2.

Ranging Backoff Start: Initial back-off window for initial ranging contention, expressed as a power of two. Values range 0 to 15 (the highest order bits are unused and set to 0).

Ranging Backoff End: Final back-off window for initial ranging contention is expressed as a power of two. Values range 0 to 15 (the highest order bits are 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 to 15 (the highest order bits are unused and set to 0). See clause 7.2.2.1.2, for an explanation of how this value is used by DOCSIS 3.0 CMs operating in Multiple Transmit Channel Mode to determine backoff on a bonding group.

Data Backoff End: Final back-off window for contention data and requests, expressed as a power of two. Values range 0 to 15 (the highest order bits are be unused and set to 0). See clause 7.2.2.1.2, for an explanation of how this value is used by DOCSIS 3.0 CMs operating in Multiple Transmit Channel Mode to determine backoff on a bonding group.

MAP Information Elements: Describe the specific usage of upstream intervals as detailed below:

The CMTS MUST comply with figure 6-19 and table 6-27 for MAP Information Elements. Values for IUCs are defined in table 6-27 and are described in detail in clause 7.2.2.1.2.

NOTE: Refer to clause 7.2.1.1, The Allocation MAP MAC Management Message, for the relationship between Alloc Start/Ack Time and the timebase.

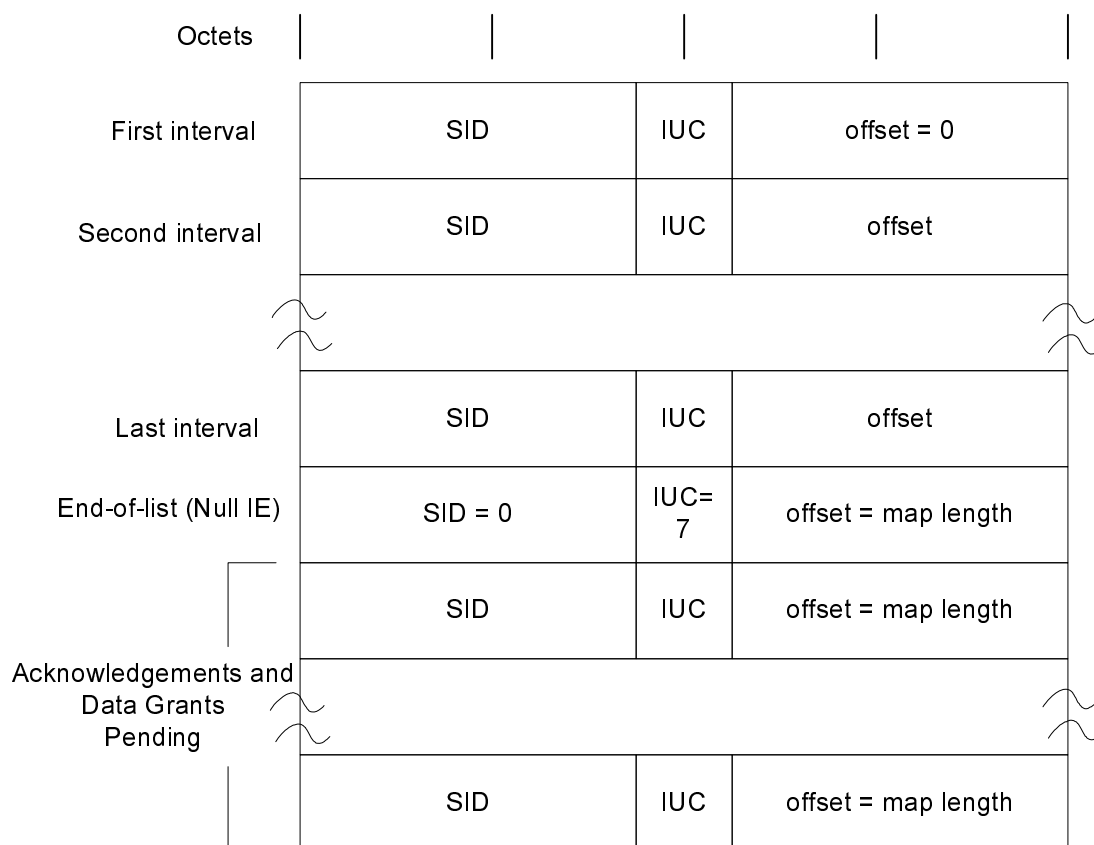


Figure 6-19: MAP Information Element Structure

Table 6-27: Allocation MAP Information Elements (IE)

IE Name (see note 1)	Interval Usage Code (IUC) (4 bits)	SID (14 bits)	Mini-slot Offset (14 bits)
Request (see note 6)	1	any	Starting offset of REQ region
REQ/Data (refer to annex A for multicast definition)	2	multicast	Starting offset of REQ/Data region (well-known multicasts define start intervals)
Initial Maintenance (see note 2)	3	broadcast or unicast	Starting offset of MAINT region (used in Initial or Periodic Ranging)
Station Maintenance	4	unicast (see note 3)	Starting offset of MAINT region (used in Periodic Ranging)
Short Data Grant (see note 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
Advanced PHY Short (see note 5) Data Grant	9	unicast	Starting offset of Data Grant assignment; if inferred length = 0, then it is a Data Grant pending
Advanced PHY Long Data Grant	10	unicast	Starting offset of Data Grant assignment; if inferred length = 0, then it is a Data Grant pending
Advanced PHY Unsolicited Grant	11	unicast	Starting offset of Data Grant assignment
Reserved	12 to 14	any	Reserved
Expansion	15	expanded IUC	# of additional 32-bit words in this IE
NOTE 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.			
NOTE 2: The CMTS MUST NOT use a unicast SID with an initial maintenance IUC on any upstream that is not a Type 3 or 4 Upstream Channel.			
NOTE 3: The SID used by the CM in the Station Maintenance IE MUST be a Temporary SID or the Ranging SID that was assigned in the REG-RSP message to the CM. For Pre-3.0 DOCSIS CMs, this MUST be the Primary SID or the Temporary SID (see clause 7.2.1.2.4).			
NOTE 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. For Multiple Transmit Channel Mode, the CM does not make any assumptions on the burst descriptor to use based on the request size and the CMTS does not necessarily grant opportunities using burst descriptors based on the amount requested or size of granted segments.			
NOTE 5: The Advanced PHY types are provided for channels carrying a combination of DOCSIS 1.x and DOCSIS 2.0/3.0 bursts and also for channels carrying DOCSIS 2.0/3.0 bursts only.			
NOTE 6: The CMTS MUST ensure that the Request IE is large enough to hold a Queue-Depth based request. Since the Queue-Depth based request and Pre-3.0 DOCSIS request frames are different sizes, the PHY parameters for IUC1 need to be carefully chosen so that the same number of mini-slots is required to hold both frame sizes.			

6.4.5 Ranging Request Messages

Ranging Request messages are transmitted by a CM at initialization on an upstream and periodically on request from the CMTS to determine network delay and request power adjustment. There are three types of Ranging Request messages: RNG-REQ, INIT-RNG-REQ and B-INIT-RNG-REQ. Table 6-28 shows when each type of message is used.

A CM MUST transmit a RNG-REQ message periodically on request from the CMTS for all unicast maintenance opportunities the CMTS provides. A CM MUST transmit a RNG-REQ at initialization on a Type 1 or 2 upstream when an MDD (MAC Domain Descriptor) is not present. A CM MUST transmit an INIT-RNG-REQ at initialization on a Type 3 or 4 upstream when an MDD is not present. A CM MUST transmit a B-INIT-RNG-REQ at initialization when it has received an MDD.

Table 6-28: CM Ranging Request Type Usage

Ranging Situation	Channel Type	
	1, 2	3, 4
CM receiving an MDD, initializing on first channel and transmitting in a broadcast Initial Maintenance opportunity	B-INIT-RNG-REQ	B-INIT-RNG-REQ
CM initializing on secondary channel and transmitting in a broadcast or unicast Initial Maintenance opportunity	INIT-RNG-REQ	INIT-RNG-REQ
CM not receiving an MDD and transmitting in a broadcast Initial Maintenance opportunity	RNG-REQ	INIT-RNG-REQ
CM transmitting in a Station Maintenance opportunity	RNG-REQ	RNG-REQ
NOTE 1: Initializing on a channel refers to the CM's first ranging request attempt during initialization and all subsequent ranging request transmissions on that channel prior to receiving a ranging response message.		
NOTE 2: First channel refers to the channel (or multiple channels in failure scenarios) on which the CM attempts to range prior to receiving the first ranging response during initialization.		
NOTE 3: Secondary channel refers to any channel on which the CM attempts to range after receiving a ranging response on a different channel, except where that ranging response contained an Upstream Channel ID Override.		

If Upstream Transmit Power Reporting is enabled in the MDD message (see clause 6.4.28.1.12), the CM MUST use the SSAP field of the MAC Management Message Header of RNG-REQ, INIT-RNG-REQ and B-INIT-RNG-REQ messages to report its Transmit Power Level, P_r , for the upstream channel on which the message is transmitted. The power level MUST be expressed by the CM in units of 1/4 dB. If there are no MDD messages present or if the Power Reporting TLV is not present in the MDD messages or if the Power Reporting is disabled by the Power Reporting TLV, then the CM MUST NOT report its Power Level in the SSAP field of the ranging messages.

If the CM has been properly commanded by the CMTS to adjust the transmitter parameters on one of its channels, it will find a Reconfiguration Time [12] in order to make the adjustment (see clause 10.3). If the CM has been properly commanded by the CMTS to adjust its dynamic range window it will wait for a Global Reconfiguration Time [12] to make the adjustment. If the CM constructs a RNG-REQ or INIT-RNG-REQ message while adjustments to the transmitter parameters for that channel (including dynamic range window adjustments) are pending, it MUST set the value of the SSAP field to 255 (Adjustment Pending) to indicate that an adjustment was pending and may or may not have been completed when the message was transmitted. The CMTS SHOULD NOT make an adjustment in the CM's transmitter parameters based on a message with a value of 255 (Adjustment Pending) in the SSAP field.

If the CM is reporting Tx Power Level or Adjustment Pending using the SSAP field it MUST set the RSVD field to zero. In this case, the CMTS MUST ignore any information in the RSVD field. If the CM detects an error condition with respect to its dynamic range window and a received RNG-RSP message, the CM MUST set bits 15 and 14 of the SID field in subsequent RNG-REQ messages until the error is cleared as follows:

Bits 15 to 14 of SID field:

00 = No error condition.

01 = Power Adjustment not applied - Commanded power adjustment would cause P_r to be outside of the 12 dB dynamic range window.

10 = The current value for P_r is more than 3dB below the top of the dynamic range window for all channels. Spurious and noise requirements are relaxed for modem operating in this condition [12].

11 = Maximum Scheduled Codes Unnecessary - MSC and Power Headroom were sent in RNG-RSP, but current P_r is sufficient to allow use of all codes. The CM does not ignore MSC setting in RNG-RSP, but just indicates a possible error condition with this encoding.

If the CM is reporting its Transmit Power in the RNG-REQ messages and Maximum Scheduled Codes (MSC) is enabled in the CMTS, the CMTS is in full control of the MSC feature. In this case, if the CMTS needs to command an increase in the Transmit Power Level which would result in the CM having a non-zero power shortfall, the CMTS MUST proactively send Maximum Scheduled Codes and Power Headroom in the RNG-RSP message.

If Upstream Transmit Power Reporting is not enabled in the MDD message, the CM MUST set the RSVD field of the MAC Management Message Header to report support for S-CDMA MSC if and only if MSC has been enabled in the UCD for this channel. In this case, the CM MUST report the maximum ratio of number of active codes to Maximum Scheduled Codes that the CM can support. The CMTS will use this value in calculating an appropriate value for Maximum Scheduled Codes to assign to the CM. The CM MUST support a Maximum Ratio of 32.

When the CM reports MSC information, the CM MUST also report its current transmit power shortfall (in dB). The CM power shortfall is the difference between the current target transmit power of the ranging request and the maximum SCDMA spreader-on transmit power of 53 dBmV. The CM MUST report a power shortfall of 0 if the current target transmit power of the ranging request is less than or equal to 53 dBmV. This value will be used by the CMTS for calculating appropriate values for S-CDMA Maximum Scheduled Codes and S-CDMA Power Headroom for the CM.

The format of the RSVD field for conveying its current transmit power shortfall when MSC is supported by the CMTS is:

Bit 7: 1= S-CDMA Maximum Scheduled Codes Supported	Bits 6 to 5: CM Maximum Ratio of: $\frac{\text{Number_Active_Codes}}{\text{Maximum_Scheduled_Codes}}$ 00 = 2 01 = 8 10 = 16 11 = 32	Bit 4 to 0: CM power shortfall (1/4 dB)
--	---	---

6.4.5.1 Ranging Request (RNG-REQ)

The RNG-REQ message transmitted by the CM MUST use an FC_TYPE = MAC Specific Header and FC_PARM = Timing MAC Header, followed by a Packet PDU in the format shown in figure 6-20.

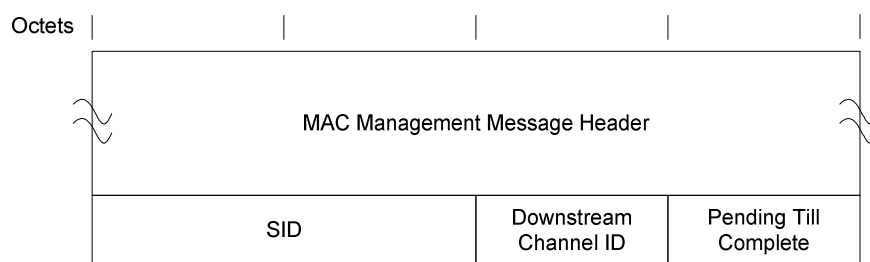


Figure 6-20: RNG-REQ Format

The parameters of RNG-REQ messages transmitted by the CM MUST be as follows.

SID: For RNG-REQ messages transmitted in Broadcast 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 downstream and upstream channels as directed by a downloaded parameter file.
- Primary SID (previously assigned in REG-RSP) for Pre-3.0 DOCSIS operation, if modem is registered and is changing upstream channels or if the CM is redoing initial ranging as a result of a DCC, UCC or UCD change (see clauses 6.4.3 and 11.1).

For RNG-REQ messages transmitted in Unicast Initial Maintenance or Station Maintenance intervals:

- Temporary SID if modem has not yet registered.
- Ranging SID if one has been assigned by the CMTS to the CM for this channel.
- Primary SID (previously assigned in REG-RSP) for Pre-3.0 DOCSIS operation, if modem is registered or is redoing initial ranging as a result of DCC, UCC or UCD change.

This is a 16-bit field of which the lower 14 bits define the SID.

Downstream Channel ID: The identifier of the downstream channel on which the CM is receiving the UCDs and MAPs which describe this upstream. This is an 8-bit field.

Pending Till Complete: If zero, then all previous Ranging Response attributes have been applied prior to transmitting this request. If non zero, this is time estimated to be needed to complete assimilation of ranging parameters. Note that only equalization can be deferred. Units are in unsigned centi-seconds (10 ms).

6.4.5.2 Initial Ranging Request (INIT-RNG-REQ)

The INIT-RNG-REQ message transmitted by the CM MUST use an FC_TYPE = MAC Specific Header and FC_PARM = Timing MAC Header, followed by a Packet PDU in the format shown in figure 6-21. The INIT-RNG-REQ differs from the RNG-REQ in that it has an upstream channel ID in place of the Pending Till Complete field in a RNG-REQ.

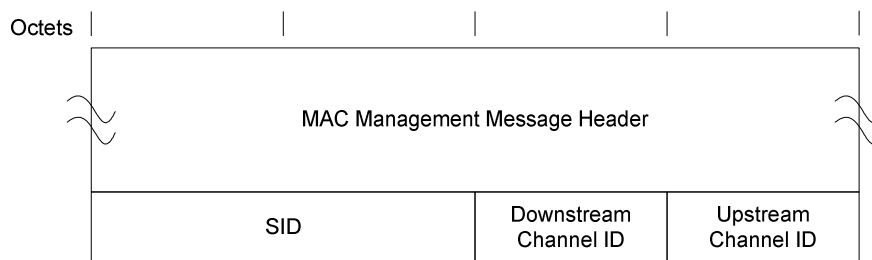


Figure 6-21: INIT-RNG-REQ Format

The parameters of the INIT-RNG-REQ message transmitted by the CM MUST be as follows.

SID:

- 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 downstream and upstream channels, as directed by a downloaded parameter file.
- Ranging SID (previously assigned in REG-RSP) if the modem is doing initial ranging as part of upstream channel acquisition.
- Ranging SID (previously assigned in REG-RSP) if modem is registered and is changing upstream channels.
- Ranging SID (previously assigned in REG-RSP) if the modem is redoing initial ranging as a result of a DBC, DCC or UCD change (see clauses 6.4.3 and 11.1).
- Ranging SID (previously assigned in REG-RSP) if the modem is redoing initial ranging as a result of a temporary loss of downstream signal, while in S-CDMA mode (see clause 10.2.2).
- Primary SID (previously assigned in REG-RSP), for CMs operating in Pre-3.0 DOCSIS mode, if modem is registered and is changing upstream channels, if the CM is redoing initial ranging as a result of a DCC, UCC or UCD change (see clauses 6.4.3 and 11.1) or if the CM is redoing initial ranging as a result of a temporary loss of downstream signal, while in S-CDMA mode (see clause 10.2.2).

This is a 16-bit field of which the lower 14 bits define the SID.

Downstream Channel ID: The identifier of the downstream channel on which the CM is receiving the UCDs and MAPs which describe this upstream. This is an 8-bit field.

Upstream Channel ID: The Upstream Channel ID from the UCD the CM is using to transmit this INIT-RNG-REQ. In the case where multiple logical upstreams are sharing the same spectrum and the Broadcast Initial Ranging Opportunities of some of these logical channels are aligned, the Upstream Channel ID allows the CMTS to know which logical channel the CM is using.

6.4.5.3 Bonded Initial Ranging Request (B-INIT-RNG-REQ)

The B-INIT-RNG-REQ message transmitted by a CM MUST use an FC_TYPE = MAC Specific Header and FC_PARM = Timing MAC Header, followed by a Packet PDU in the format shown in figure 6-22.

The B-INIT-RNG-REQ differs from the INIT-RNG-REQ in that it includes the MD-DS-SG-ID used for downstream topology resolution and a set of Capability Flags in place of the SID. A CM MUST only use this message for the first channel it ranges on when an MDD is present. When ranging for the first time on all succeeding channels, the CM uses the INIT-RNG-REQ message (see clause 6.4.5.2).




Figure 6-22: B-INIT-RNG-REQ Format

The parameters of the B-INIT-RNG-REQ message transmitted by the CM MUST be as follows.

Capability Flags: Used to convey modem capabilities that are needed prior to registration by the CMTS. It is an 8-bit field as defined in clause 6.4.5.3.1.

MD-DS-SG-ID: The identifier of the MAC Domain Downstream Service Group obtained from downstream ambiguity resolution. This is an 8-bit field. The value zero indicates that the MD-DS-SG-ID could not be determined.

Downstream Channel ID: The identifier of the downstream channel on which the CM is receiving the UCDs and MAPs which describe this upstream. This is an 8-bit field.

Upstream Channel ID: The Upstream Channel ID from the UCD the CM is using to transmit this B-INIT-RNG-REQ. In the case where multiple logical upstreams are sharing the same spectrum and the Broadcast Initial Ranging Opportunities of some of these logical channels are aligned, the Upstream Channel ID allows the CMTS to know which logical channel the CM is using.

If the MD-DS-SG-ID is unrecognized, the CMTS MUST silently ignore the B-INIT-RNG-REQ.

6.4.5.3.1 Capability Flags

A CM MUST indicate certain capabilities to the CMTS prior to registration via the Capability Flags field. The format of the Capability Flags field used by the CM MUST be as follows.

Table 6-29: Capability Flags Encoding

Bit 7: 1: Pre-3.0 DOCSIS fragmentation is supported prior to registration 0: Pre-3.0 DOCSIS fragmentation is not supported prior to registration	Bit 6: 1: Early Authentication and Encryption Supported 0: Early Authentication and Encryption Not Supported	Bits 5 to 0: Reserved
--	--	--------------------------

A CM MUST indicate support for pre-3.0 DOCSIS fragmentation prior to registration.

A CM MUST indicate support for Early Authentication and Encryption.

6.4.6 Ranging Response (RNG-RSP)

A Ranging Response MUST be transmitted by a CMTS in response to received RNG-REQ, INIT-RNG-REQ or B-INIT-RNG-REQ. The state machines describing the ranging procedure appear in clause 10.2.3.7. 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 by the CMTS in a type/length/value (TLV) form.

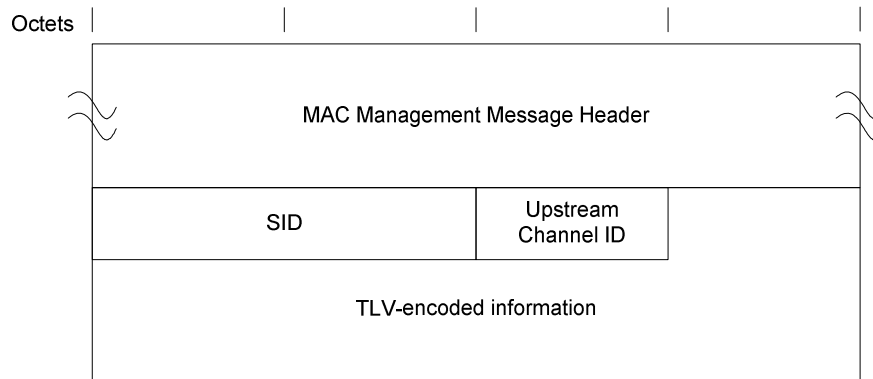


Figure 6-23: Ranging Response

A CMTS MUST generate Ranging Responses in the form shown in figure 6-23, including all of the following parameters as defined below:

SID: If the modem is being instructed by this response to move to a different channel, this is the initialization SID. If this is a response to an initial ranging request (whether RNG-REQ, INIT-RNG-REQ or B-INIT-RNG-REQ), this is the assigned temporary SID. Otherwise, this is the SID from the corresponding RNG-REQ to which this response refers.

Upstream Channel ID: The identifier of the upstream channel on which the CMTS received the RNG-REQ, INIT-RNG-REQ or B-INIT-RNG-REQ to which this response refers. On the first ranging response received by the CM after initializing or reinitializing it is MAC, this channel ID may be different from the channel ID the CM used to transmit the range request. Thus, the CM MUST use this channel ID for the rest of its transactions, not the channel ID from which it initiated the range request.

All other parameters, when present, MUST be coded as TLV tuples and used by the CMTS, as defined below:

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

Timing Adjust, Integer Part: The amount by which to change the Ranging Offset of the burst transmission so that bursts arrive at the expected mini-slot time at the CMTS. The units are $(1 / 10,24 \text{ MHz}) = 97,65625 \text{ ns}$. A negative value implies the Ranging Offset is to be decreased, resulting in later times of transmission at the CM (see clause 6.4.20, table 6-10 and clause 6.2).

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 clause 6.4.6.4.)

Upstream Channel ID Override: An optional parameter. The identifier of the upstream channel with which the modem should redo initial ranging. (See clause 6.4.6.4.)

Timing Adjust, Fractional Part: Higher resolution timing adjust offset to be appended to Timing Adjust, Integer Part. The units are $(1 / (256 * 10,24 \text{ MHz})) = 0,38 \text{ ns}$. This parameter provides finer granularity timing offset information. This TLV is a mandatory parameter for timing adjustments on S-CDMA channels. This TLV is an optional parameter for timing adjustments on TDMA channels. A CM whose timing is locked to the downstream symbol clock **MUST** apply the fractional part timing adjustment if this TLV is present, whether the channel is TDMA or S-CDMA.

S-CDMA Maximum Scheduled Codes: The value that the CMTS uses to limit the number of codes scheduled to a CM in an S-CDMA frame. CMs that implement the S-CDMA Maximum Scheduled Codes use this value to limit the maximum size of a concatenated burst in an S-CDMA Frame.

S-CDMA Power Headroom: CMs that implement the S-CDMA Maximum Scheduled Codes use this value to control transmit power as per [12] when Maximum Scheduled Codes is Enabled.

Upstream Channel Adjustments: A CMTS can send this TLV to move a CM to another upstream channel as a part of upstream ambiguity resolution and to adjust more than one upstream channel with a single RNG-RSP message when a modem has Multiple Transmit Channel Mode enabled.

T4 Timeout Multiplier: A CMTS can send this TLV to increase the value of the T4 timeout for CMs that have Multiple Transmit Channel Mode enabled.

6.4.6.1 Encodings

The type values used by the CMTS in the RNG-RSP **MUST** comply with table 6-30 and figure 6-24. These are unique within the ranging response message but not across the entire MAC message set. The type and length fields used by the CMTS in the RNG-RSP **MUST** each be 1 octet in length.

Table 6-30: Ranging Response Message Encodings

Name	Type (1 byte)	Length (1 byte)	Value (Variable Length)
Timing Adjust, Integer Part	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 to be convolved with current values (refer to [12]). The CMTS MUST NOT include this TLV in a RNG-RSP that includes a type 9 TLV.
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.
Timing Adjust, Fractional Part	8	1	TX timing fine offset adjustment. 8-bit unsigned value specifying the fine timing adjustment in units of $1 / (256 * 10,24 \text{ MHz})$.
Transmit Equalization Set	9	n	TX equalization data to be loaded in place of current values (refer to [12]). The CMTS MUST NOT include this TLV in a RNG-RSP to a DOCSIS 1.x CM. The CMTS MUST NOT include this TLV in a RNG-RSP that includes a type 4 TLV.
S-CDMA Maximum Scheduled Codes	10	1	A CMTS may send this TLV only if a CM indicated that it supports the S-CDMA Maximum Scheduled Codes. A value of 0 means no code limit. Other possible values range from 4 to number_active_codes inclusive. Maximum Scheduled codes is an integer multiple of codes_per_mini-slot. The CMTS MUST NOT include this TLV if S-CDMA mode is disabled. Absence of this TLV indicates that Maximum Scheduled Codes is inactive for this CM, which MUST then use the S-CDMA Number of Active Codes.
S-CDMA Power Headroom	11	1	A CMTS sends this TLV to a CM in conjunction with TLV-10. The CMTS MUST NOT include this TLV if S-CDMA mode is disabled. The units are dB. The range of this TLV is from 0 to $4 * 10 \log \left(\frac{\text{Number_Active_Codecs}}{\text{Maximum_Scheduled_Codecs}} \right)$ (see note).

Name	Type (1 byte)	Length (1 byte)	Value (Variable Length)
Upstream Channel Adjustments	12	n	A CMTS may send one or more sets of this TLV to allow for adjustments to channels other than the one provided in the RNG-RSP message or for use in ambiguity resolution.
Upstream Channel ID	12.1	1	The ID of the channel.
Temp SID	12.2	2	SID to be used on the new channel.
Initialization Technique	12.3	1	1 = Perform broadcast initial ranging. 2 = Perform unicast ranging. 3 = Perform either broadcast or unicast ranging. 0, 4 to 255: reserved.
Ranging Parameters	12.4	n	Contains sub-TLVs for ranging adjustments.
Deprecated	12.4.1	1	
Timing Offset, Integer Part	12.4.2	4	TX timing offset adjustment (signed 32-bit, units of (6,25 microsec/64)).
Timing Offset, Fractional Part	12.4.3	1	TX timing fine offset adjustment. 8-bit unsigned value specifying the fine timing adjustment in units of $1 / (256 * 10,24 \text{ MHz})$.
Power Offset	12.4.4	1	TX Power offset adjustment (signed 8-bit, 1/4-dB units).
Frequency Offset	12.4.5	2	TX frequency offset adjustment (signed 16-bit, Hz units).
Ranging Status	12.4.6		1 = continue, 2 = abort, 3 = success. The Ranging Status sub-TLV is not applicable during US Ambiguity Initial Ranging and is only used after Registration.
T4 Timeout Multiplier	13	1	Multiplier of the default T4 Timeout as defined earlier in this clause. If omitted the default as defined in annex B is used. The valid range is 1 to 10.
Dynamic Range Window Upper Edge	14	1	The upper edge of the Dynamic Range Window expressed in units 1/4 db below the max allowable setting (Phi) [12]. The CM does not need this value prior to registration and an equivalent TLV is provided in the TCC encodings so that the CMTS can communicate the setting to the CM during registration.
Reserved	15 to 255	n	Reserved for future use.

NOTE: A value of 0 for TLV-10 restricts the range to 0 for TLV-11.

6.4.6.2 Example of TLV Data

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

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

Figure 6-24: Example of TLV Encoded Data

6.4.6.3 Transmit Equalization Encodings

Type 4 or 9	Length	Main Tap Location	Number of Forward Taps per Symbol
Number of Forward Taps (N)	Reserved		
First Coefficient F_1 (real)		First Coefficient F_1 (imag)	
\approx			
Last Coefficient F_N (real)		Last Coefficient F_N (imag)	

Figure 6-25: Generalized Decision Feedback Equalization Coefficients

The number of taps per modulation interval T signaled by the CMTS 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 T-spaced equalizer, the number of taps per modulation interval field MUST be set to "1" by the CMTS. The total number of taps signaled by the CMTS 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 or 9 elements may be used. Data MUST be treated by the CM and CMTS as if byte-concatenated, that is, the first byte after the length field of the second type 4 or 9 element is treated as if it immediately followed the last byte of the first type 4 or 9 element.

6.4.6.4 RNG-RSP Channel Overrides

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. After transmitting the first RNG-RSP with Ranging Status equal to Success(3) to an initializing CM, the CMTS MUST NOT send the CM an upstream or downstream channel override in a RNG-RSP message. If a downstream frequency override is specified in the RNG-RSP, the modem MUST reinitialize its MAC (see clause 10.2.1) using initial ranging with the specified downstream center frequency as the first scanned channel.

If an upstream channel ID override is specified in the RNG-RSP, the modem MUST reinitialize its MAC (see clause 10.2.1) 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 (refer to clause 10.2.1) 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 de-assigned. The modem MUST redo initial ranging using the Initialization SID.

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

6.4.6.5 Upstream Channel Adjustments

A CMTS sends this TLV for use in upstream ambiguity resolution and for post registration ranging adjustments to one or more upstream channels other than the upstream channel indicated by the Upstream Channel ID encoded in the body of the RNG-RSP message prior to the TLV encodings.

During upstream ambiguity resolution, the CMTS MUST include no more than one Upstream Channel Adjustment TLV. The CMTS MUST include the Upstream Channel ID and Initialization Technique sub-TLVs. The CMTS MUST include the Temp SID TLV when the Initialization Technique includes "unicast ranging" (techniques 2 and 3). The CMTS MUST NOT include the Temp SID TLV when the Initialization Technique is "broadcast initial ranging" (technique 1). The CMTS MUST NOT send a Downstream Frequency Override TLV when an Upstream Channel Adjustment TLV is present. The CMTS MUST NOT send an Upstream Channel ID Override TLV when an Upstream Channel Adjustment TLV is present. The CMTS MAY send Ranging Parameter sub-TLVs to speed up upstream ambiguity resolution. During upstream ambiguity resolution, the CMTS MUST NOT include Ranging Response parameter adjustments (adjustments specified in TLVs 1 through 4 and 8 through 11) in the RNG-RSP containing the Upstream Channel Adjustment TLV. The CMTS MUST NOT include the Ranging Status sub-TLV in the Upstream Channel Adjustment TLVs during upstream ambiguity resolution.

The CMTS MUST NOT include this TLV in a RNG-RSP message between the completion of US Ambiguity Initial Ranging and receiving a REG-ACK. During this period the CM will only have a single US channel and should not be moved to other US channels via this method.

After registration, the CMTS MAY include one or more Upstream Channel Adjustment TLVs in a RNG-RSP message during periodic station maintenance to adjust multiple US channels with a single RNG-RSP message. In this case, the Temp SID and Initialization Technique sub-TLVs MUST NOT be present. The Upstream Channel ID field will represent the UCID of the channel to be adjusted and the Ranging Parameters field will indicate what adjustments are to be made to that channel. The presence of Upstream Channel Adjustments for a particular upstream channel will reset the T3 timer, if active, for that channel. If the CMTS does not include the Ranging Status sub-TLV in the Upstream Channel Adjustment TLVs, the CM MUST consider the Ranging Status to be unchanged.

Prior to the completion of US Ambiguity Initial Ranging, the CM MUST change US channels in accordance with the parameters in this TLV. If the Ranging Parameter sub-TLVs (12.4.1 through 12.4.6) are used, the CM MUST apply the offsets as referenced to the current values for the channel on which the RNG-REQ message was sent. After completion of US Ambiguity Initial Ranging and prior to sending the REG-ACK, the CM MUST ignore an Upstream Channel Adjustment TLV if present in a RNG-RSP. After sending the REG-ACK, the CM assumes that the Upstream Channel Adjustments TLV indicates changes to be made to upstream channels other than the upstream channel indicated in the body of the RNG-RSP message prior to the TLV encodings and MUST adjust transmissions on those upstream channels according to the TLV parameters. As a result, if the CM receives US Channel Adjustments with an unknown US Channel ID after registration, it MUST ignore that TLV.

6.4.6.6 T4 Timeout Multiplier

In Multiple Transmit Channel Mode the CMTS MAY increase the value of the T4 timeout by means of the T4 Timeout Multiplier in order to reduce CMTS overhead associated with scheduling RNG-REQ slots and processing RNG-RSP messages. The CM MUST set its T4 timeout to the value of the multiplier times the default T4 timeout in annex B. If a RNG-RSP does not contain a T4 Timeout Multiplier value then the CM MUST use the default T4 timeout as defined in annex B. If the CMTS includes a T4 Timeout Multiplier in the RNG-RSP, the CMTS MUST set it to be in the valid range of 1 to 10. In order to allow for future updates, the CM does not enforce the valid range.

If the CMTS sets the T4 Timeout Multiplier to any value other than the default then it MUST send the T4 Timeout Multiplier value in every RNG-RSP. When reducing the value of the T4 Timeout Multiplier, the CMTS SHOULD start scheduling a few RNG-REQ slots at the shorter interval before sending a RNG-RSP with the shorter timeout. When increasing the value of the T4 Timeout Multiplier, the CMTS SHOULD continue scheduling a few RNG-REQ slots at the shorter interval even after the RNG-RSP with the longer value is transmitted.

6.4.7 Registration Request Messages

The CM transmits a Registration Request message after receipt of a CM configuration file as specified in clause 10.2. There are two types of Registration Request message in DOCSIS 3.0: the single frame Registration Request message (referred to as REG-REQ) and the Multipart Registration Request message (referred to as REG-REQ-MP). If the Primary Downstream Channel upon which the CM is registering does not contain a valid MDD message, then the CM MUST transmit a (REG-REQ) message. If the Primary Downstream Channel upon which the CM is registering does contain a valid MDD message, then the CM MUST transmit a REG-REQ-MP message instead of a REG-REQ. This specification will use the terms "REG-REQ" and "REG-REQ-MP" when it is important to make a distinction between the two and the term "Registration Request" when such a distinction is not necessary.

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 CM MUST include the following Configuration Settings in the Registration Request:

Configuration File Settings:

- All configuration settings included in the pre-3.0 DOCSIS CMTS MIC calculation as specified in clause D.2.1.
- All TLVs selected by the E-MIC Bitmap (if the Extended CMTS MIC Encoding TLV is present in the configuration file).
- The following "allowed unprotected" TLVs:
 - Enable 2.0 Mode.
 - Downstream Channel List.
 - CMTS MIC Configuration Setting.
 - Channel Assignment Configuration Settings.
 - Upstream Drop Classifier Group ID.

NOTE: The CM MUST forward DOCSIS Extension Field 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 CM MUST NOT include the Configuration Settings not in the above list in the Registration Request message.

The CM MUST include the Vendor ID Configuration Setting (Vendor ID of CM) registration parameters in the Registration Request.

The CM MUST include the Modem Capabilities Encodings registration parameter in the Registration Request. The CM MUST specify all of its Modem Capabilities in its Registration Request, subject to the restrictions in clause 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.

The CM MUST include one or more Receive Channel Profile Encodings registration parameters in the REG-REQ-MP.

The CM MAY include the following registration parameters in the Registration Request:

- Modem IP Address.
- Vendor Specific Capabilities.

The Vendor Specific Capabilities field is for vendor specific information not included in the configuration file.

6.4.7.1 Registration Request (REG-REQ)

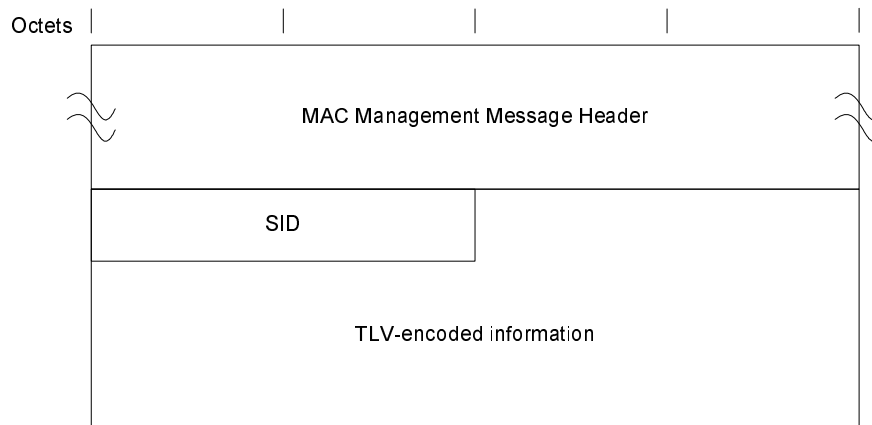


Figure 6-26: Registration Request (REG-REQ)

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 annex C.

If the configuration file causes the CM to create a REG-REQ larger than 1 518 bytes then the CM MUST retry downloading the configuration file up to TFTP Request Retries as defined in annex B. If the CM is unable to download the requested file from the TFTP server that will result in a REG-REQ message of acceptable size, then the CM MUST declare IP Connectivity failed. The CM MUST log an event in the local log server indicating the cause of the failure.

6.4.7.2 Multipart Registration Request (REG-REQ-MP)

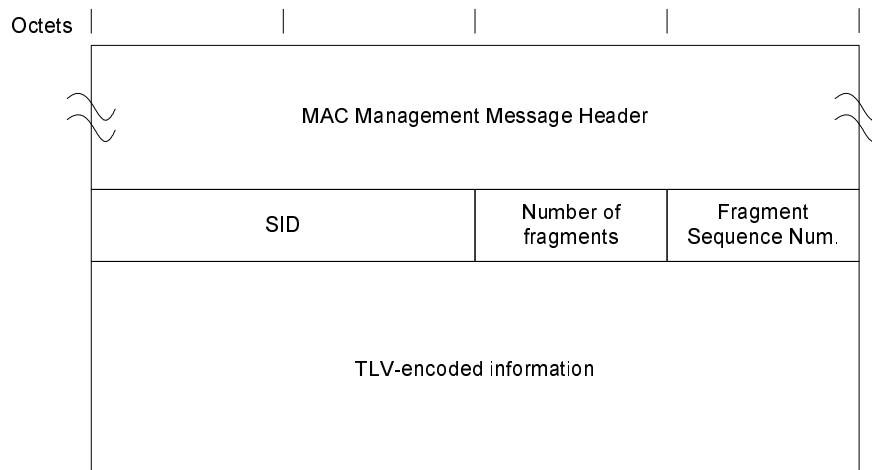


Figure 6-27: Multipart Registration Request (REG-REQ-MP)

A CM MUST generate Multipart Registration Requests in the form shown in figure 6-27, including the following parameters:

SID: Temporary SID for this CM.

Number of Fragments: Fragmentation allows the REG-REQ-MP TLV parameters to be spread across more than one DOCSIS MAC Frame, thus allowing the total size of the REG-REQ-MP to exceed the maximum payload of a single MAC management frame. The value of this field represents the number of REG-REQ-MP MAC management frames that a unique and complete set of REG-REQ-MP TLV parameters are spread across to constitute the complete REG-REQ-MP message. This field is an 8-bit unsigned integer.

Fragment Sequence Number: This field indicates the position of this fragment in the sequence that constitutes the complete REG-REQ-MP message. Fragment Sequence Numbers start with the value of 1 and increase by 1 for each fragment in the sequence. Thus, the first REG-REQ-MP message fragment has a Fragment Sequence Number of 1 and the last REG-REQ-MP message fragment has a Fragment Sequence Number equal to the Number of Fragments. The CM MUST NOT fragment any top level TLVs of a REG-REQ-MP. Each REG-REQ-MP message fragment is a complete DOCSIS frame with its own CRC. Other than the Fragment Sequence Number, the framing of one REG-REQ-MP message fragment is independent of the framing of another REG-REQ-MP message fragment. This potentially allows the CMTS to process fragments as they are received rather than reassembling the entire payload. This field is an 8-bit unsigned integer.

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

The CMTS MUST be capable of receiving a REG-REQ-MP containing a total MAC Management payload size of at least 16 000 bytes.

The MAC Management Message Type value, Number of Fragments field and Fragment Sequence Number field distinguish the REG-REQ-MP from the REG-REQ. In all other respects, the REG-REQ-MP is identical to the REG-REQ (clause 6.4.7.1).

6.4.8 Registration Response Messages

There are two types of Registration Response message in DOCSIS 3.0: the single frame Registration Response message (referred to as REG-RSP) and the Multipart Registration Response message (referred to as REG-RSP-MP). This specification will use the terms "REG-RSP" and "REG-RSP-MP" when it is important to make a distinction between the two and the term "Registration Response" when such a distinction is not necessary.

The CMTS transmits a Registration Response or Multipart Registration Response after receipt of a CM Registration Request or Multipart Registration Request (respectively).

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

Service Flow Parameters: All the Service Flow Parameters from the REG-REQ or REG-REQ-MP, plus the Service Flow ID assigned by the CMTS. Every Service Flow that contained a Service Class Name that was admitted/activated, is expanded into the full set of TLVs defining the Service Flow. Every upstream Service Flow that was admitted/activated, has 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 or REG-REQ-MP, plus the assigned Service Flow ID.

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

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

If the REG-REQ or REG-REQ-MP 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 annex C, the CMTS MUST format the REG-RSP or REG-RSP-MP to contain at least one of the following:

Service Flow Error Set: A Service Flow Error Set and identifying Service Flow Reference is included for at least one failed Service Flow in the corresponding REG-REQ or REG-REQ-MP. Every Service Flow Error Set includes at least one specific failed QoS Parameter of the corresponding Service Flow.

Classifier Error Set: A Classifier Error Set and identifying Classifier Reference and Service Flow Reference is included for at least one failed Classifier in the corresponding REG-REQ or REG-REQ-MP. Every Classifier Error Set includes at least one specific failed Classifier Parameter of the corresponding Classifier.

Payload Header Suppression Error Set: A PHS Error Set and identifying Service Flow Reference and Classifier Reference pair is included for at least one failed PHS Rule in the corresponding REG-REQ or REG-REQ-MP. Every PHS Error Set includes 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 REG-REQ or REG-REQ-MP contains a Service Flow Reference and a Service Class Name for deferred admission/activation, the CMTS MUST NOT include any additional QoS Parameters except the Service Flow Identifier in the REG-RSP or REG-RSP-MP (refer to clause 7.5.3).

If the corresponding REG-REQ or REG-REQ-MP contains DOCSIS 1.0 Service Class TLVs (refer to clause C.1.1.4), the CMTS MUST format the REG-RSP or REG-RSP-MP to contain the following TLV tuples:

DOCSIS 1.0 Service Class Data: Returned when Response = Okay. This is a Service ID / service class tuple for each class of service granted.

NOTE: Service class IDs included by the CMTS MUST be those requested in the corresponding REG-REQ or REG-REQ-MP.

Service Not Available: Returned when Response = Class of Service Failure. If a service class cannot be supported, this configuration setting is returned in place of the service class data.

If the CMTS is returning a non-zero value for the Multiple Transmit Channel Support modem capability encoding to put the modem into a Multiple Transmit Channel Mode of operation, the REG-RSP or REG-RSP-MP MUST include:

- The Transmit Channel Configuration.
- The Service Flow SID Cluster Assignment.

If the CMTS is returning a non-zero value for the Multiple Receive Channel Support modem capability encoding to put the modem into a Multiple Receive Channel mode of operation, the REG-RSP or REG-RSP-MP MUST include:

- The Receive Channel Configuration.
- DSID Encodings.

All other parameters are coded TLV tuples:

Security Association Encodings: In certain cases a REG-RSP or REG-RSP-MP transmitted by a CMTS can also contain Security Association Encodings (refer to clauses 9.2.3, 9.2.4 and C.1.5.5).

Modem Capabilities: The CMTS response to the capabilities of the modem.

Vendor-Specific Data: As defined in clause C.1.1.18.2.

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

6.4.8.1 Registration Response (REG-RSP)

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

To provide for flexibility, the message parameters following the Response field **MUST** be encoded by the CMTS in a TLV format.

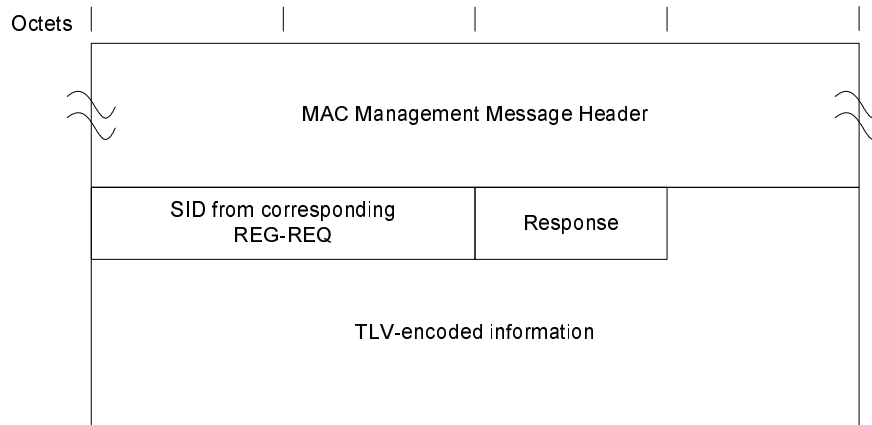


Figure 6-28: Registration Response Format

A CMTS **MUST** generate Registration Responses in the form shown in figure 6-28, 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 with DOCSIS 1.0 Class of Service Encodings:

0 = Okay.

1 = Authentication Failure.

2 = Class of Service Failure.

For REG-RSP to a modem registering with Service Flow Encodings, this field contains one of the Confirmation Codes in clause C.4.

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.

6.4.8.2 Multipart Registration Response (REG-RSP-MP)

A Multipart Registration Response **MUST** be transmitted by the CMTS in response to a received REG-REQ-MP.

To provide for flexibility, the message parameters following the Response field **MUST** be encoded by the CMTS in a TLV format.

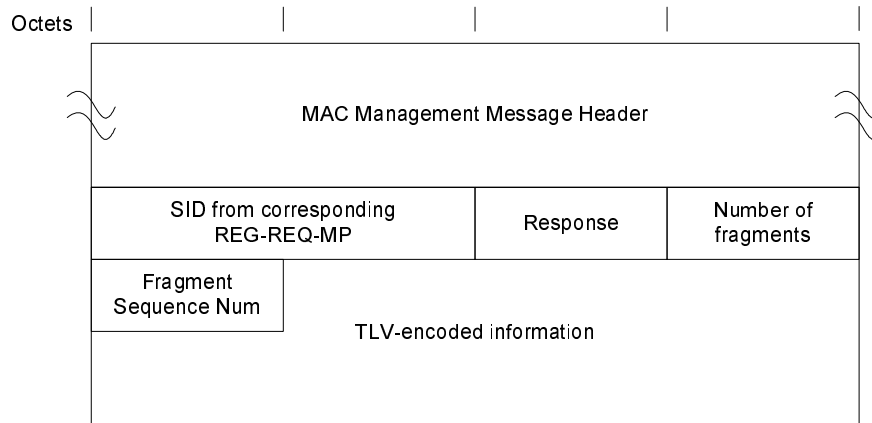


Figure 6-29: Multipart Registration Response Format

A CMTS **MUST** generate Multipart Registration Responses in the form shown in figure 6-29, including the following parameters:

SID from Corresponding REG-REQ-MP: SID from corresponding REG-REQ-MP to which this response refers (this acts as a transaction identifier).

Response:

For REG-RSP-MP to a modem registering with DOCSIS 1.0 Class of Service Encodings:

0 = Okay.

1 = Authentication Failure.

2 = Class of Service Failure.

For REG-RSP-MP to a modem registering with Service Flow Encodings, this field contains one of the Confirmation Codes in clause C.4.

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.

Number of fragments: Fragmentation allows the REG-RSP-MP TLV parameters to be spread across more than one DOCSIS MAC Frame, thus allowing the total size of the REG-RSP-MP to exceed the maximum payload of a single MAC management frame. The value of this field represents the number of REG-RSP-MP MAC management frames that a unique and complete set of REG-RSP-MP TLV parameters are spread across to constitute the REG-RSP-MP message. This field is an 8-bit unsigned integer. The number of fragments in the REG-RSP-MP can differ from the number of fragments in the REG-REQ-MP to which this response refers.

Fragment Sequence Number: This field indicates the position of this fragment in the sequence that constitutes the complete REG-RSP-MP message. Fragment Sequence Numbers start with the value of 1 and increase by 1 for each fragment in the sequence. Thus, the first REG-RSP-MP message fragment has a Fragment Sequence Number of 1 and the last REG-RSP-MP message fragment has a Fragment Sequence Number equal to the Number of Fragments. The CMTS **MUST** send the message fragments in order of increasing sequence numbers. The CMTS **MUST NOT** fragment any top level TLVs across message fragments of a REG-RSP-MP. Each REG-RSP-MP message fragment is a complete DOCSIS frame with its own CRC. Other than the Fragment Sequence Number, the framing of one REG-RSP-MP message fragment is independent of the framing of another REG-RSP-MP message fragment. This potentially allows the CM to process fragments as they are received rather than reassembling the entire payload. This field is an 8-bit unsigned integer.

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

The CM MUST be capable of receiving a REG-RSP-MP containing a total MAC Management payload size of at least 16 000 bytes.

The MAC Management Message Type value, Number of Fragments field and Fragment Sequence Number field distinguish the REG-RSP-MP from the REG-RSP. In all other respects, the REG-RSP-MP is identical to the REG-RSP (clause 6.4.8.1).

6.4.8.3 Encodings

The type values used by the CMTS 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 used by the CMTS MUST each be 1 octet.

6.4.8.3.1 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 is setting a capability to "on" (indicating that it may be used), unless explicitly indicated otherwise in the subclauses of clause C.1.3.1, the CMTS MUST return the capability TLV to the CM with the same value as the CM included in the Registration Request. If the CMTS does not recognize a modem capability, it MUST return the TLV with the value zero ("off") in the Registration Response. The CMTS MUST NOT include a capability in the Registration Response that was not present in the corresponding Registration Request.

Only capabilities set to "on" in the Registration Request may be set "on" in the Registration Response as this is the handshake indicating that they have been successfully negotiated. Capabilities set to "off" in the Registration Request MUST also be set to "off" in the Registration Response by the CMTS.

Encodings are as defined for the Registration Request.

6.4.8.3.2 DOCSIS 1.0 Service Class Data

The CMTS MUST include a DOCSIS 1.0 Service Class Data parameter in the Registration Response for each DOCSIS 1.0 Class of Service parameter (refer to clause C.1.1.4) in the Registration Request.

This encoding defines the parameters associated with a requested class of service. 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 by the CMTS to define the parameters for a single service class. The CMTS MUST use multiple class definitions for multiple service class data configuration setting sets.

If each received DOCSIS 1.0 Class of Service parameter does not have a unique Class ID in the range 1..16, the CMTS MUST send a REG-RSP or REG-RSP-MP with a class-of-service failure response and no DOCSIS 1.0 Class-of-Service TLVs. If no Class ID is present for any single DOCSIS 1.0 Class-of-Service TLV in the REG-REQ or REG-REQ-MP, the CMTS MUST send a REG-RSP or REG-RSP-MP 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 specifies the identifier for the class of service to which the encapsulated string applies. The CMTS MUST set this to be the Class ID which was requested in the associated REG-REQ or REG-REQ-MP if present.

Type	Length	Value
1.1	1	Class ID from REG-REQ or REG-REQ-MP

Service ID: The CMTS MUST set the value of the field to specify the SID associated with this service class.

Type	Length	Value
1.2	2	SID

6.4.9 Registration Acknowledge (REG-ACK)

A Registration Acknowledge MUST be transmitted by the CM in response to a REG-RSP or REG-RSP-MP from the CMTS under the circumstances described in clause 10.2.6.1. It confirms acceptance by the CM of the Registration Response parameters as reported by the CMTS. The CM MUST format a REG-ACK as shown in figure 6-30.

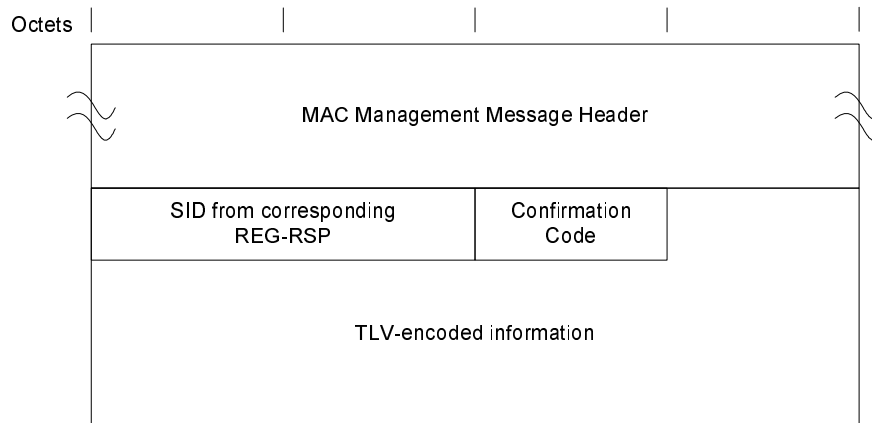


Figure 6-30: Registration Acknowledgment

The parameters of the REG-ACK transmitted by the CM 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 clause 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 Registration Request (see clause 6.4.7). The CMTS will return them with Identifiers, expanding Service Class Names if present, in the Registration Response (see clause 6.4.8). Since the CM may be unable to support one or more of these provisioned items, the Registration Acknowledge defines Error Sets for all failures related to these provisioned items.

If there were any failures of provisioned items, the CM MUST include in the REG-ACK the Error Sets corresponding to those failures as described below. The Error Set identification is provided by using Service Flow ID and Classifier ID from corresponding REG-RSP or REG-RSP-MP. If a Classifier ID or SFID was omitted in the REG-RSP or REG-RSP-MP, the CM MUST use the appropriate Reference (Classifier Reference, SF Reference) in the REG-ACK.

Classifier Error Set: A Classifier Error Set and identifying item is included for at least one failed Classifier in the corresponding configuration file, REG-RSP or REG-RSP-MP. For QoS Classifiers, the identifying item is the Classifier Reference/Identifier and Service Flow Reference/Identifier pair and the failed Classifier occurs in the corresponding REG-RSP or REG-RSP-MP message. For Upstream Drop Classifiers, the identifying item is the Classifier Identifier and the failed classifier occurs in either the configuration file, REG-RSP or REG-RSP-MP message, depending on the location of the Upstream Drop Classifiers that the CM uses for filtering (clause 7.5.1.2.2). Every Classifier Error Set includes at least one specific failed Classifier Parameter of the corresponding Classifier. This parameter is omitted if the entire Registration Request/Response 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 or REG-RSP-MP message. A Service Flow Error Set and identifying Service Flow Reference/Identifier is included for at least one failed QoS Parameter of at least one failed Service Flow in the corresponding REG-RSP or REG-RSP-MP message. This parameter is omitted if the entire Registration Request/Response is successful.

Payload Header Suppression Error Set: A PHS Error Set and identifying Service flow Reference/Identifier and Classifier Reference/Identifier pair is included for at least one failed PHS Rule in the corresponding REG-RSP or REG-RSP-MP. Every PHS Error Set includes at least one specific failed PHS of the failed PHS Rule. This parameter is omitted if the entire Registration Request/Response is successful.

TCC Error Set: A TCC Error Set and identifying TCC Reference is included for at least one failed TCC in the corresponding REG-RSP. Every TCC Error Set includes at least one specific failed parameter of the corresponding TCC. It does not need to include every failed parameter of the corresponding TCC. This parameter is omitted if the entire Registration Request/Response is successful (refer to clause C.1.5.1).

RCC Error Set: An RCC Error Set is included to report an error in an RCC encoding in the corresponding REG-RSP. Every RCC Error Set includes at least one specific failed parameter of the corresponding RCC. It does not need to include every failed parameter of the corresponding RCC. This parameter is omitted if the entire Registration Request/Response is successful (refer to clause C.1.5.3).

In the case where the CM is unable to acquire one or more of the upstream and/or downstream channels assigned via the TCC and/or RCC encodings (respectively), the CM needs to report back to the CMTS the list of channels that it was unable to acquire so that the CMTS can take appropriate action. If the CM is unable to acquire one or more of the downstream channels assigned to it in the RCC, the CM MUST include an RCC encoding with a Partial Service Downstream Channels TLV in the REG-ACK, which includes a list of the downstream channels that could not be acquired. If the CM is unable to acquire one or more of the upstream channels assigned to it in the TCC, the CM MUST include a TCC encoding with a TCC Error Encoding for each upstream channel it was unable to acquire in the REG-ACK, corresponding to the TCC encoding that assigned that upstream channel in the REG-RSP. This is because each TCC encoding describes the actions to take for a single upstream channel. Note that this is different from the case of reporting an error in the encoding, where only a single error needs to be reported (even if multiple errors exist).

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 clause 7.5.3). 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 send a REG-ACK with error sets if it has insufficient resources to actually support this Service Flow.

6.4.10 Upstream Channel Change Request (UCC-REQ)

An Upstream Channel Change Request MAY be transmitted by a CMTS to cause a DOCSIS 1.0 CM to change the upstream channel on which it is transmitting. This message has been superseded by the DCC-REQ message (introduced in DOCSIS RFI 1.1) and is only present for backwards compatibility. A CMTS MUST NOT transmit a UCC-REQ to a DOCSIS 1.1 or later CM. However, for backward compatibility, a CM MUST support the receipt of an UCC-REQ message.

The format of an UCC-REQ message transmitted by the CMTS MUST be as shown in figure 6-31.

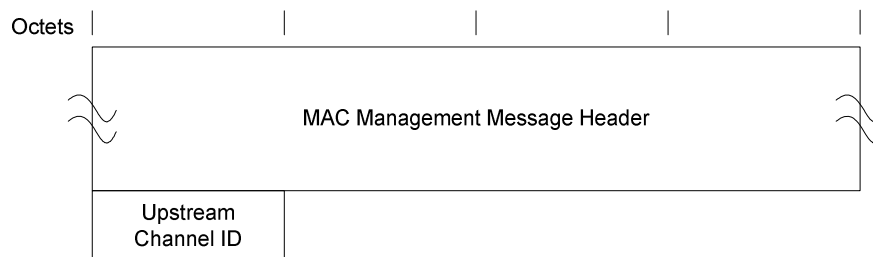


Figure 6-31: Upstream Channel Change Request

The parameters of a UCC-REQ transmitted by the CMTS 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.

6.4.11 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 an UCC-RSP message transmitted by the CMTS **MUST** be as shown in figure 6-32.

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 proceed without re-performing registration. The full procedure for changing channels is described in clause 11.3.

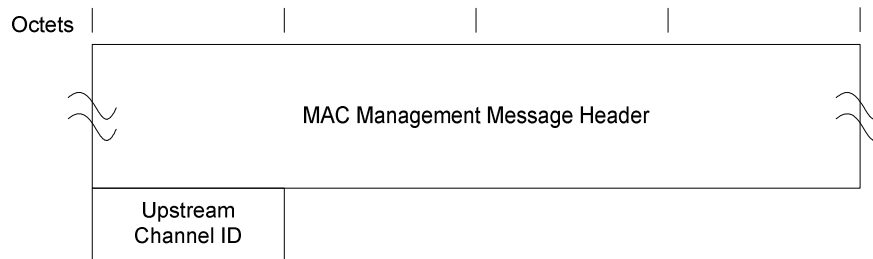


Figure 6-32: Upstream Channel Change Response

The parameters of an UCC-RSP transmitted by a CM **MUST** be as follows:

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

6.4.12 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.

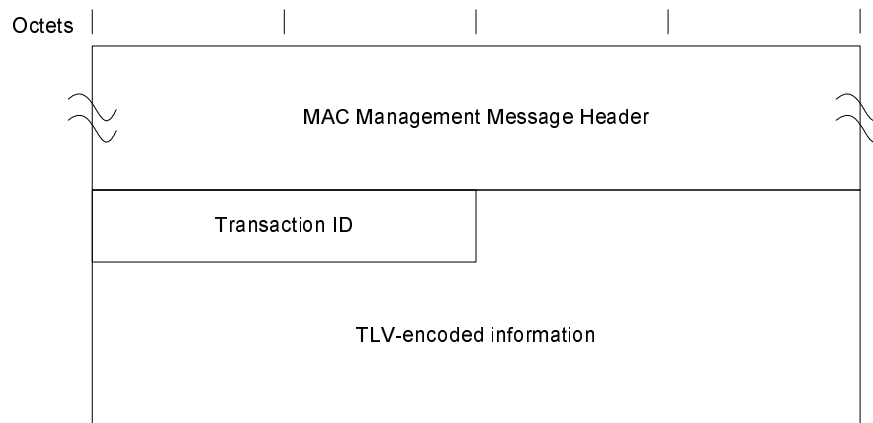


Figure 6-33: Dynamic Service Addition - Request

A CM or CMTS **MUST** generate DSA-REQ messages in the form shown in figure 6-33 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 annex C. A DSA-REQ message transmitted by a CM or CMTS **MUST NOT** contain parameters for more than one Service Flow in each direction, i.e. a DSA-REQ message contains 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 transmitted by a CM or CMTS MUST contain:

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

The DSA-REQ message transmitted by a CM or CMTS MAY contain classifier parameters and payload header suppression parameters associated with the Service Flows specified in the message. If included, the CM or CMTS MUST comply with the following rules for classifier parameters and payload header suppression parameters:

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 transmitted by a CM or CMTS MUST contain:

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the Dynamic Service message's Attribute list (refer to clause C.1.4.1).

6.4.12.1 CM-Initiated Dynamic Service Addition

The CM MUST use the Service Flow Reference to link Classifiers to Service Flows when generating a CM initiated DSA-REQ. 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 by the CM. This value need not be unique with respect to the other service flows known by the sender.

The CM MUST use the Classifier Reference and Service Flow Reference to link Payload Header Suppression Parameters to Classifiers and Service Flows when generating a CM-initiated DSA-REQ. Values of the Classifier Reference are local to the DSA message; each Classifier within the DSA-request MUST be assigned a unique Classifier Reference by the CM.

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

6.4.12.2 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.

CMTS-initiated DSA-Requests sent to a CM in a Multiple Transmit Channel Mode of operation MUST include Service Flow SID Cluster Assignments.

6.4.13 Dynamic Service Addition - Response (DSA-RSP)

A Dynamic Service Addition Response **MUST** be generated in response to a received DSA-Request by a CM or CMTS. The format of a DSA-RSP used by a CM or CMTS **MUST** be as shown in figure 6-34.

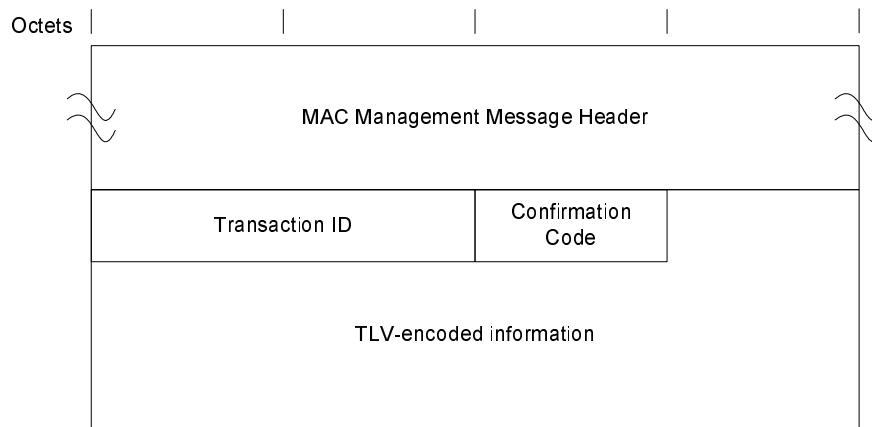


Figure 6-34: Dynamic Service Addition - Response

The parameters of DSA-RSP transmitted by a CM or CMTS **MUST** be as follows:

Transaction ID: Transaction ID from corresponding DSA-REQ.

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

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

If the transaction is successful, the DSA-RSP contains 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 clause C.4, the DSA-RSP transmitted by the CM or CMTS **MUST** contain at least one of the following:

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

If Privacy is enabled, the DSA-RSP message transmitted by the CM or CMTS MUST contain:

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the Dynamic Service message's Attribute list (refer to clause C.1.4.1).

6.4.13.1 CM-Initiated Dynamic Service Addition

The CMTS's DSA-Response for Service Flows that are successfully added MUST contain a Service Flow ID. The CMTS's 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 clause C.2.2.3.4) to request service addition, the CMTS's 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 CM received TCC Encodings in the Registration Response, the CMTS MUST include Service Flow SID Cluster Assignments.

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.

6.4.13.2 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.

6.4.14 Dynamic Service Addition - Acknowledge (DSA-ACK)

A Dynamic Service Addition Acknowledge MUST be generated by a CM or CMTS in response to a received DSA-RSP. The format of a DSA-ACK transmitted by a CM or CMTS MUST be as shown in figure 6-35.

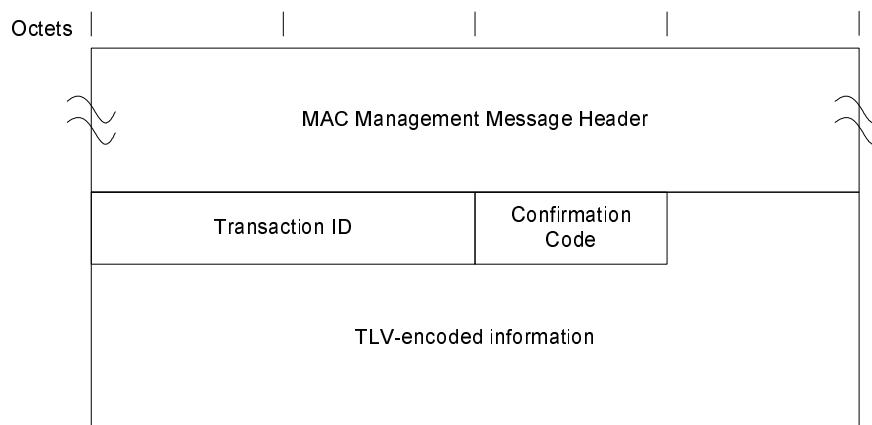


Figure 6-35: Dynamic Service Addition - Acknowledge

The parameters of a DSA-ACK transmitted by a CM or CMTS MUST be as follows:

Transaction ID: Transaction ID from corresponding DSA-Response.

Confirmation Code: The appropriate Confirmation Code (refer to clause C.4) for the entire corresponding DSA-Response.

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 is included for at least one failed QoS Parameter of at least one failed Service Flow in the corresponding DSA-REQ. This parameter is omitted if the entire DSA-REQ is successful.

If Privacy is enabled, the DSA-RSP message transmitted by the CM or CMTS MUST contain:

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the Dynamic Service message's Attribute list (refer to clause C.1.4.1).

6.4.15 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. If a CMTS sends a DSC-REQ message changing an Upstream Drop Classifier, then conceptually the Upstream Drop Classifier is associated with a NULL Service Flow that is not signaled in the DSC-REQ message. DSCs transmitted by a CM or CMTS that are changing classifiers MUST carry the entire classifier TLV set for that new classifier.

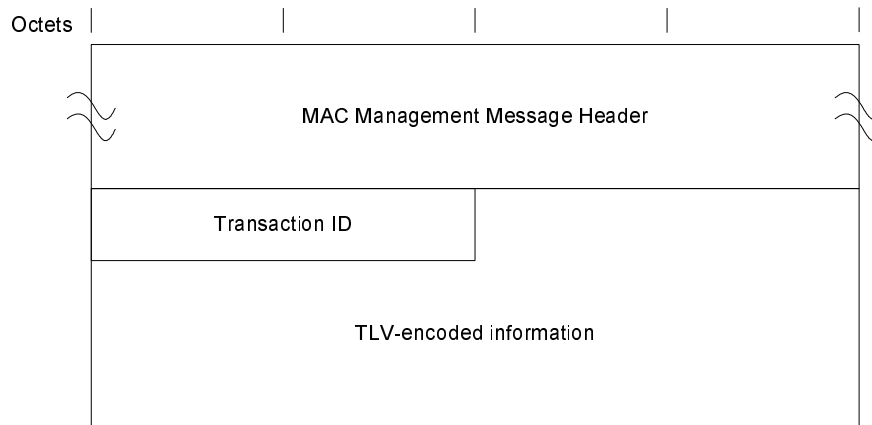


Figure 6-36: Dynamic Service Change - Request

A CM or CMTS MUST generate DSC-REQ messages in the form shown in figure 6-36 including the following parameters as described below:

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

All other parameters are coded as TLV tuples as defined in annex C. A DSC-REQ message transmitted by a CM or CMTS MUST NOT carry parameters for more than one Service Flow in each direction, i.e. a DSC-REQ message contains 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 transmitted by a CM or CMTS MUST contain at least one of the following:

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) are set to null. If Service Flow Parameters are included, they contain a Service Flow Identifier.

Not all Service Flow Parameters are permitted to be changed via a DSC-REQ message. Reference values and Identifiers (TLVs 24/25.1-3) are unique to a Service Flow and as such can not be modified by a CM or CMTS in a DSC-REQ. In addition, the following Service Flow Parameter TLVs **MUST NOT** be modified by a CM or CMTS via a DSC-REQ:

- Service Flow Scheduling Type (TLV 24.15).
- Bit 9 (Segment Header on/off) of the Request/Transmission Policy (TLV 24.16).
- Multiplier to Number of Bytes Requested (TLV 24.26).

Support for changes to the following Service Flow Parameter TLVs via a DSC-REQ is optional in a receiving CM or CMTS:

- Service Class Name (TLV 24/25.4) (only if all the parameters that differ in the new class are allowed to change).
- Service Flow Required Attribute Mask (TLV 24/25.31).
- Service Flow Forbidden Attribute Mask (TLV 24/25.32).
- Service Flow Attribute Aggregation Rule Mask (TLV 24/25.33).
- Application Identifier (TLV 24/25.34).
- Vendor Specific QoS Parameters (TLV 24/25.43).

If changes to these parameters are specified in a DSC-REQ, the receiving CM or CMTS **MAY** implement the change. Since support for these changes is optional, they might be rejected by the receiving entity. Changes to all other Service Flow Parameters via a DSC-REQ message **MUST** be supported by both CMs and CMTSs.

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 clause C.2.1.4.7). If included, the Classifier Parameters contains the Dynamic Change Action TLV, a Classifier Reference/Identifier and a Service Flow Identifier.

Not all Classifier Parameters are permitted to be changed via a DSC-REQ message. Reference values and Identifiers (TLVs 22/23/60.1-4) are unique to a Classifier and as such can not be modified by a CM or CMTS in a DSC-REQ. If changes are specified to Vendor Specific QoS Parameters (TLV 22/23/60.43) in a DSC-REQ, the receiving CM or CMTS **MAY** implement the change. Since support for changes to these parameters is optional, they might be rejected by the receiving entity. Changes to all other Classifier Parameters via a DSC-REQ message **MUST** be supported by both CMs and CMTSs.

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 clause C.2.2.8.5). If included, the PHS parameters 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". If the Dynamic Service Change Action is "Delete all PHS Rules", the PHS Parameters 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 are ignored by the receiver of the DSC-REQ message.

Not all Payload Header Suppression Parameters are permitted to be changed via a DSC-REQ message. Reference values and Identifiers (TLVs 26.1-4) are unique to a Payload Header Suppression Rule and as such can not be modified by a CM or CMTS in a DSC-REQ. If changes are specified to Vendor Specific QoS Parameters (TLV 26.43) in a DSC-REQ, the receiving CM or CMTS **MAY** implement the change. Since support for changes to these parameters is optional, they might be rejected by the receiving entity. Changes to Add and Delete PHS elements via a DSC-REQ **MUST** be supported by receiving CMs and CMTSs. Changes to Set PHS elements via a DSC-REQ **MUST** be supported by the receiving CM or CMTS when the rule is partially defined. Changes to Set PHS elements **MUST** be rejected by the receiving CM or CMTS once the rule is fully defined.

If Privacy is enabled, a DSC-REQ transmitted by the CM or CMTS **MUST** also contain:

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the Dynamic Service message's Attribute list (refer to clause C.1.4.1).

6.4.16 Dynamic Service Change - Response (DSC-RSP)

A Dynamic Service Change Response **MUST** be generated by a CM or CMTS in response to a received DSC-REQ.

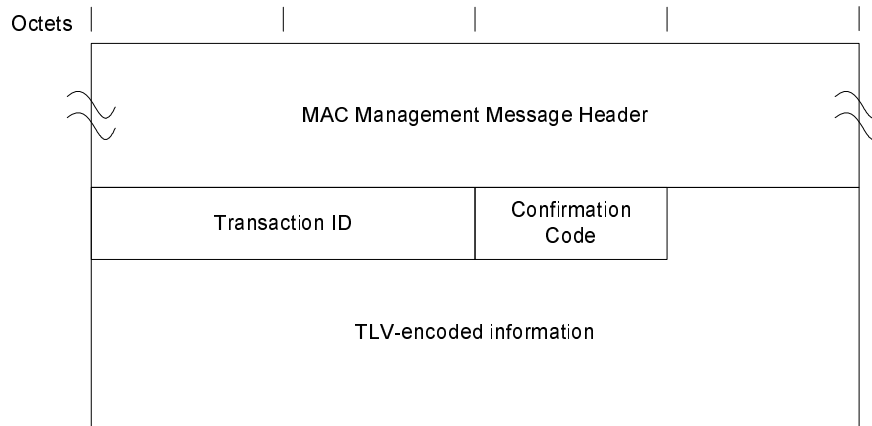


Figure 6-37: Dynamic Service Change - Response

A CM or CMTS **MUST** generate DSC-RSP messages in the form shown in figure 6-37 including the following parameters as described below:

Transaction ID: Transaction ID from the corresponding DSC-REQ.

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

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

If the transaction is successful, the DSC-RSP contains 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 Service Flow request which also included a Service Class Name, 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 annex C, the DSC-RSP transmitted by the CM or CMTS **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 is included for at least one failed Classifier in the corresponding DSC-REQ. Every Classifier Error Set includes at least one specific failed Classifier Parameter of the corresponding Classifier. This parameter is omitted if the entire DSC-REQ is successful.

Service Flow Error Set: A Service Flow Error Set and identifying Service Flow ID is included for at least one failed Service Flow in the corresponding DSC-REQ. Every Service Flow Error Set includes at least one specific failed QoS Parameter of the corresponding Service Flow. This parameter is 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 are 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) include an identifying Service Flow ID. Every PHS Error Set includes at least one specific failed PHS Parameter of the corresponding failed PHS Rule. This parameter is omitted if the entire DSC-REQ is successful".

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

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the Dynamic Service message's Attribute list (refer to clause C.1.4.1).

6.4.17 Dynamic Service Change - Acknowledge (DSC-ACK)

A Dynamic Service Change Acknowledge MUST be generated by a CM or CMTS in response to a received DSC-RSP.

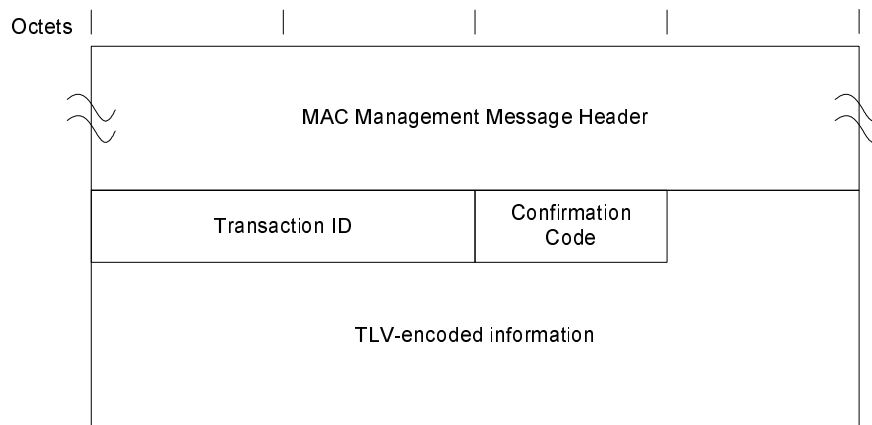


Figure 6-38: Dynamic Service Change - Acknowledge

A CM or CMTS MUST generate DSC-ACK messages in the form shown in figure 6-38 including the following parameters as described below:

Transaction ID: Transaction ID from the corresponding DSC-REQ.

Confirmation Code: The appropriate Confirmation Code (refer to clause C.4) for the entire corresponding DSC-Response.

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 is included for at least one failed QoS Parameter of at least one failed Service Flow in the corresponding DSC-REQ. This parameter is omitted if the entire DSC-REQ is successful.

If Privacy is enabled, the DSC-ACK message transmitted by the CM or CMTS MUST contain:

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the Dynamic Service message's Attribute list (refer to clause C.1.4.1).

6.4.18 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.

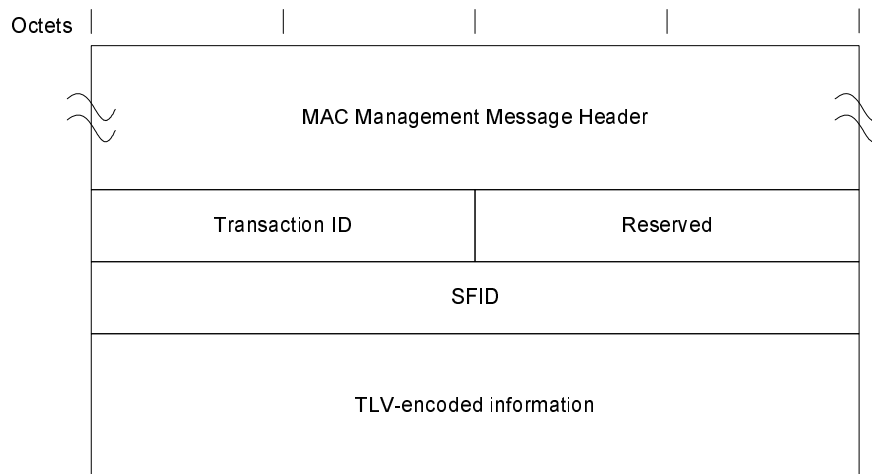


Figure 6-39: Dynamic Service Deletion - Request

A CM or CMTS MUST generate DSD-REQ messages in the form shown in figure 6-39 including the following parameters as described below:

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 TLVs. If this value is non-zero, any SFIDs included in the TLVs are ignored.

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

Reserved: Used to align the message along 32-bit boundaries.

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

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

If Privacy is enabled, the DSD-REQ transmitted by a CM or CMTS MUST include:

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the Dynamic Service message's Attribute list (refer to clause C.1.4.1).

6.4.19 Dynamic Service Deletion - Response (DSD-RSP)

A DSD-RSP MUST be generated by a CM or CMTS in response to a received DSD-REQ.

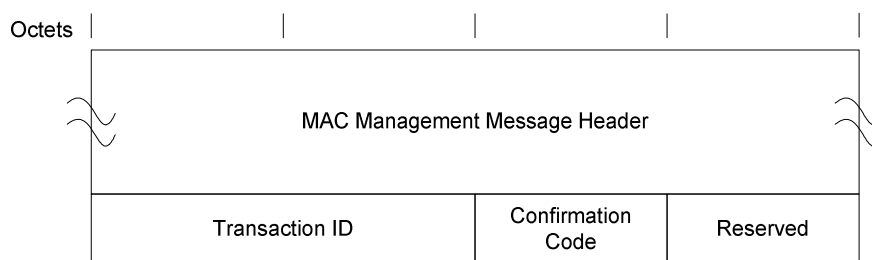


Figure 6-40: Dynamic Service Deletion - Response

A CM or CMTS MUST generate DSD-RSP messages in the form shown in figure 6-40 including the following parameters as described below:

Transaction ID: Transaction ID from the corresponding DSD-REQ.

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

Reserved: Used to align the message along 32-bit boundaries.

6.4.20 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 on which it is receiving or both. The CMTS MUST support the ability to generate DCC-REQ messages.

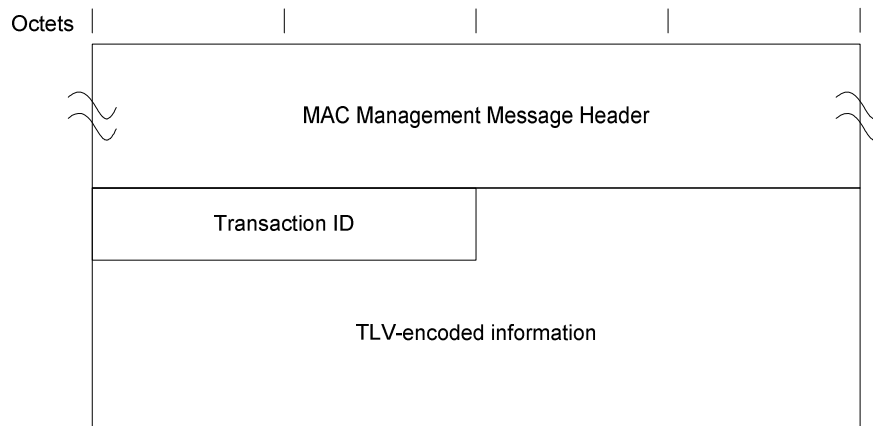


Figure 6-41: Dynamic Channel Change Request

A CMTS MUST generate DCC-REQ messages in the form shown in figure 6-41 including the following parameter:

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

The following parameters 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 and other related parameters 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. This TLV occurs as many times as necessary to contain one UCD.

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 and Service Identifier. This TLV is repeated for every Service Flow which has parameters requiring substitution.

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

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the Dynamic Channel Change message's Attribute list (refer to clause C.1.4.1).

6.4.20.1 Encodings

The type values used by the CMTS in a DCC-REQ MUST be those shown below. These are unique within the Dynamic Channel Change Request message, but not across the entire MAC message set.

6.4.20.1.1 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 SHOULD 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 by the CMTS if the upstream channel is changed, even if the UCD substitution TLV is included.

Type	Length	Value
1	1	0 to 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.

6.4.20.1.2 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	List of subtypes

The CMTS MUST include this TLV when specifying a downstream channel change. If this TLV is missing, the CM MUST NOT change its downstream parameters.

6.4.20.1.2.1 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 included by the CMTS 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.

6.4.20.1.2.2 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 to 255: reserved

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

6.4.20.1.2.3 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/s
		1 = 5,360537 Msym/s
		2 = 6,952 Msym/s
		3 to 255: reserved

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

6.4.20.1.2.4 Downstream Interleaver Depth

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

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

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

6.4.20.1.2.5 Downstream Channel Identifier

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

Subtype	Length	Value
2.5	1	0 to 255: Downstream Channel ID

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

6.4.20.1.2.6 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 not to 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 to 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 annex B) or longer if the SYNC message is lost or is not synchronized with the old channel(s).

The CM MUST observe this TLV.

6.4.20.1.3 Initialization Technique

When present, this TLV allows the CMTS to direct the CM as to what level of reinitialization, if any, the CM 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 reinitialization 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 unicast initial ranging whereas a DCC-REQ from one upstream channel group to another might require broadcast initial ranging. 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 broadcast initial ranging on new channel before normal operation
		2 = Perform unicast ranging on new channel before normal operation
		3 = Perform either broadcast or unicast ranging on new channel before normal operation
		4 = Use the new channel(s) directly without reinitializing or ranging
		5 to 255: reserved

The CMTS MUST use initialization technique 0 (re-initialize the MAC) when changing the downstream channel of a CM operating in Multiple Receive Channel Mode. The CMTS MUST use initialization technique 0 (re-initialize the MAC) when changing the upstream channel of a CM to which a Transmit Channel Configuration was assigned during registration. The CM MUST first select the new upstream and downstream channels based upon the Upstream Channel ID TLV (refer to clause 6.4.20.1) and the Downstream Frequency TLV (refer to clause 6.4.20.1.2.1). 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 clause 6.4.20.1.6.2) may specify a new primary SID for use on the new channel (refer to clause 6.4.20.1.6.2).

- Option 0: This option directs the CM to perform all the operations associated with initializing the CM (refer to clause 10.2, Cable Modem Initialization and ReInitialization). This includes all the events after acquiring downstream QAM, FEC and MPEG lock and before Standard Operation (refer to clause 7.5.7, General Operation), 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 broadcast initial ranging is specified, operation on the new channel could be delayed by several Ranging Intervals (see annex B).
- Option 2: If unicast ranging is specified, operation on the new channel could be delayed by the value of T4 (see annex B).
- Option 3: This value authorizes a CM to use an initial maintenance or station maintenance region, whichever 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. The CM may continue its normal operation as soon as it has achieved synchronization on the new channel.

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

6.4.20.1.4 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 clause 6.4.3 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 SHOULD ensure that the Upstream Channel ID for the new channel is different than the Upstream Channel ID for the old channel. If the Upstream Channel IDs for the old and new channels are identical, the CMTS MUST include this TLV. The Ranging Required parameter in the new UCD does not apply in this context, since the functionality is covered by the Initialization Technique TLV.

If the length of the UCD exceeds 254 bytes, the UCD MUST be fragmented by the CMTS into two or more successive Type 4 elements. Each fragment generated by the CMTS, 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 annex B) or longer if the UCD message is lost.

The CM MUST observe this TLV, even if the Upstream Channel ID and the UCD change count match the old channel.

6.4.20.1.5 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 CMTS MUST ensure that the baseline privacy keys associated with the SAID 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-bit 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.

6.4.20.1.6 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 indices. 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.

6.4.20.1.6.1 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 clause C.2.2.3.2 for usage details.

This TLV **MUST** be included in the DCC-REQ by the CMTS 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 CM **MUST** observe this Sub-TLV.

6.4.20.1.6.2 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 clause C.2.2.3.3 for usage details.

Subtype	Length	Value
7.2	4	Current SID (lower-order 14 bits of a 16-bit field), new SID (lower-order 14 bits of a 16-bit 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.

6.4.20.1.6.3 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 clause C.2.2.6.10 for usage details.

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 clause 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.

6.4.20.1.7 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.

Type	Length	Value
8	6	MAC Address of Destination CMTS

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

6.4.21 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 transmitted by a CM MUST be as shown in figure 6-42.

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 reinitialized 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 by the CM if it reinitializes its MAC.

The full procedure for changing channels is described in clause 11.4.

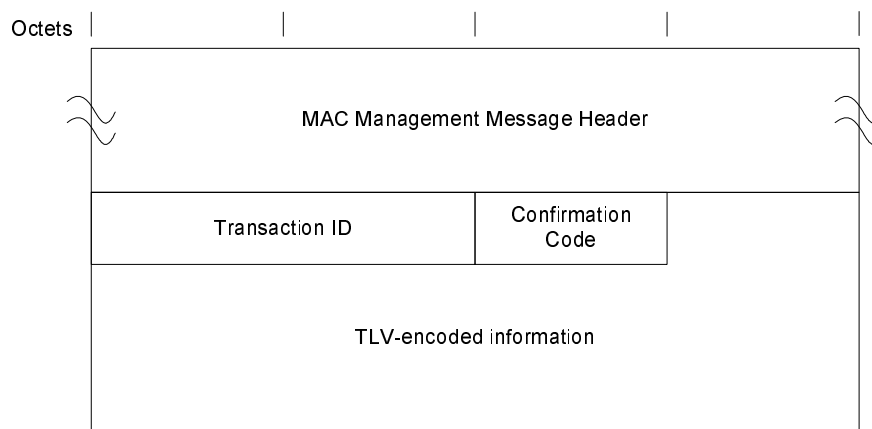


Figure 6-42: Dynamic Channel Change Response

The parameters of a DCC-RSP transmitted by a CM MUST be as follows:

Transaction ID: A 16-bit Transaction ID from the corresponding DCC-REQ.

Confirmation Code: An 8-bit Confirmation Code as described in annex C.

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 CM MUST include in the DCC-RSP:

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the Dynamic Channel Change message's Attribute list (refer to clause C.1.4.1).

6.4.21.1 Encodings

The type values used by the CM in a DCC-RSP MUST be those shown below. These are unique within the Dynamic Channel Change Response message, but not across the entire MAC message set.

6.4.21.1.1 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	List of subtypes

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.

6.4.21.1.1.1 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
1.1	4	Length (based upon timestamp)

The CM MUST include this sub-TLV if the CM Jump Time TLV is included in the DCC-RSP.

6.4.21.1.1.2 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
1.2	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 occurred. 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 if the CM Jump Time TLV is included in the DCC-RSP.

6.4.22 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 transmitted by a CMTS MUST be as shown in figure 6-43.

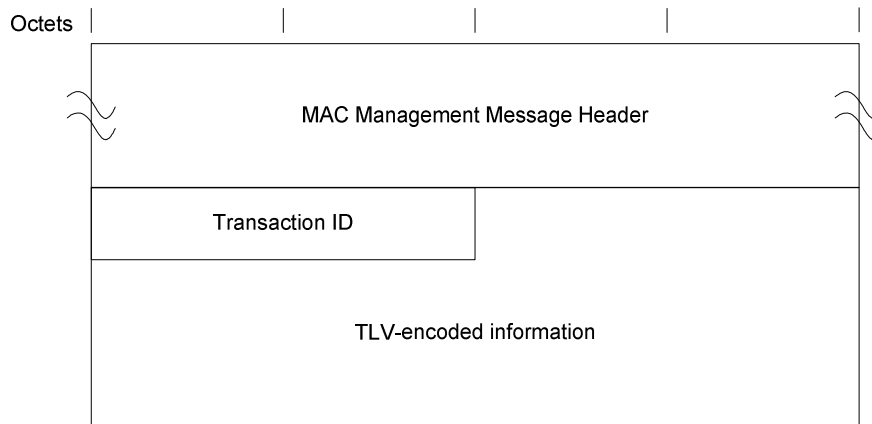


Figure 6-43: Dynamic Channel Change Acknowledge

The parameters of a DCC-ACK transmitted by a CMTS MUST be as follows:

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

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

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the Dynamic Channel Change message's Attribute list (refer to clause C.1.4.1).

6.4.23 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 annex B.

The DCI-REQ transmitted by a CM MUST be formatted as shown in figure 6-44.

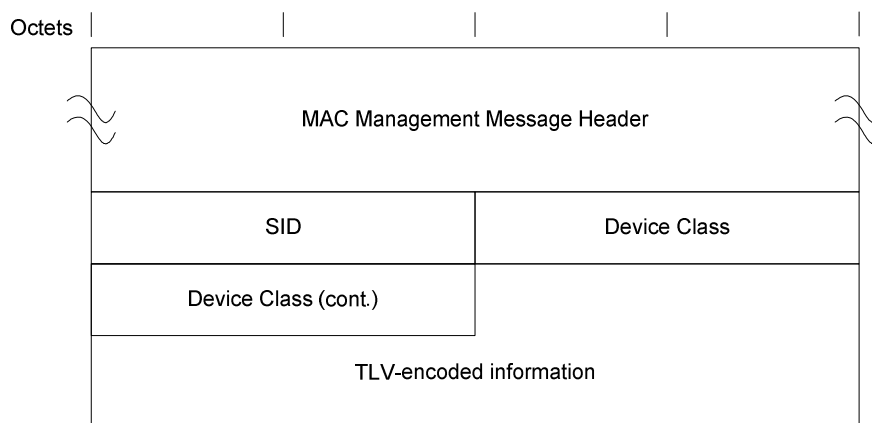


Figure 6-44: Device Class Identification Request

The parameters of a DCI-REQ transmitted by a CM 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 set to zero.

6.4.24 Device Class Identification Response (DCI-RSP)

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

The DCI-RSP transmitted by the CMTS MUST be formatted as shown in figure 6-45.

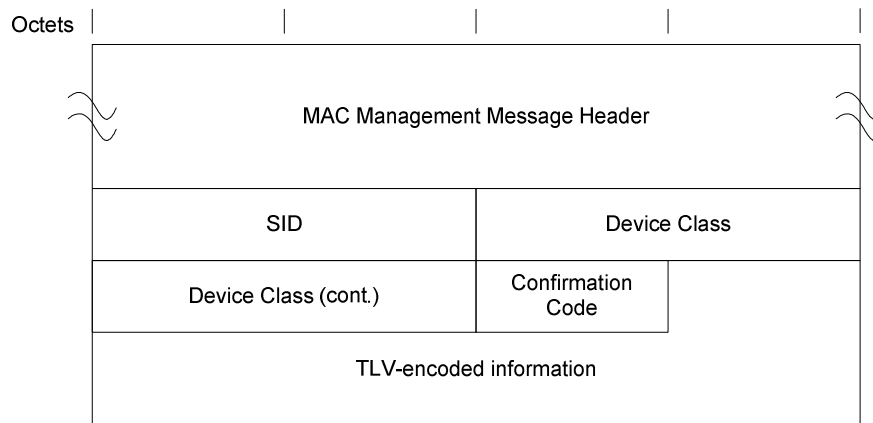


Figure 6-45: Device Class Identification Response

The parameters of a DCI-RSP transmitted by a CMTS 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 clause C.4.

The CMTS MUST use only one of 3 confirmation codes in the DCI-RSP as described below:

- 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 [50] and insert this tuple in the DHCPDISCOVER from the corresponding CM before forwarding that DHCPDISCOVER to the DHCP server.

6.4.25 Upstream Transmitter Disable (UP-DIS)

The UP-DIS message provides functionality to permanently, temporarily or for a specified period of time disable a CM's upstream transmitter(s). 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 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 annex A for details on well-known addresses).

The CMTS MAY be capable of transmitting the UP-DIS message to a CM. The CMTS can transmit UP-DIS messages either as a result of a triggering event detected by the 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, signalling, which remotely instructs the CMTS to transmit a UP-DIS message, is outside the scope of the present document. One of the possible implementations may be SNMP commands sent to the CMTS over the network.

CMs MAY support the UP-DIS message in order to ease network management.

Since the UP-DIS mechanism at the CM is stateless and the CMs do not retain disabled status after a power cycle, the CMTS MAY incorporate mechanisms to track disabled CMs by their MAC addresses. The 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 in an attempt to re-register. However, the same function may also be implemented by the provisioning infrastructure on modem registration; therefore, if the CMTS is unable to track disabled modems autonomously, it SHOULD be able to send a UP-DIS in response to an external command.

The UP-DIS message transmitted by a CMTS MUST be formatted as shown in figure 6-46.

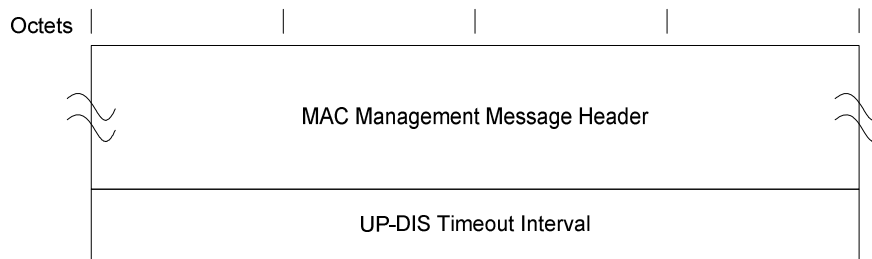


Figure 6-46: UP-DIS message format

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

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

- 00000000 disables the upstream of the modem without a timeout as described below.
- FFFFFFFF remotely reinitializes the MAC, which resumes the normal operation of the modem.

The CM MUST autonomously disable all of its upstream transmitters immediately upon receipt of an UP-DIS message with UP-DIS Timeout Interval = 0, regardless of any other transaction state (refer to clause 10) 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 disabled, the CM upstream transmitters MUST only be re-enabled by power cycling the CM or by an UP-DIS message with UP-DIS Timeout Interval = FFFFFFFF. All other UP-DIS messages MUST be ignored by the CM when the upstreams are disabled.

If supported, the CM MUST autonomously reset its upstream transmitters upon receipt of an UP-DIS message with UP-DIS Timeout Interval = FFFFFFFF, regardless of any other transaction state (refer to clause 10) 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 by the CM. If supported, the CM MUST autonomously disable its upstream transmitters 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 clause 10) or the state of its control program. Although the timeout T is specified in milliseconds, the CM MAY extend the specified time out by up to 100 ms. When the timeout expires, the CM SHOULD reinitialize the 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 by the CM, except for an UP-DIS message with UP-DIS Timeout Interval = FFFFFFFF or 00000000.

6.4.26 Test Request (TST-REQ)

The Test Request is used to force a CM to enter or leave one of two test modes. The TST-REQ message with Mode != 0 MUST NOT be sent by the CMTS except in response to an explicit command from the operator.

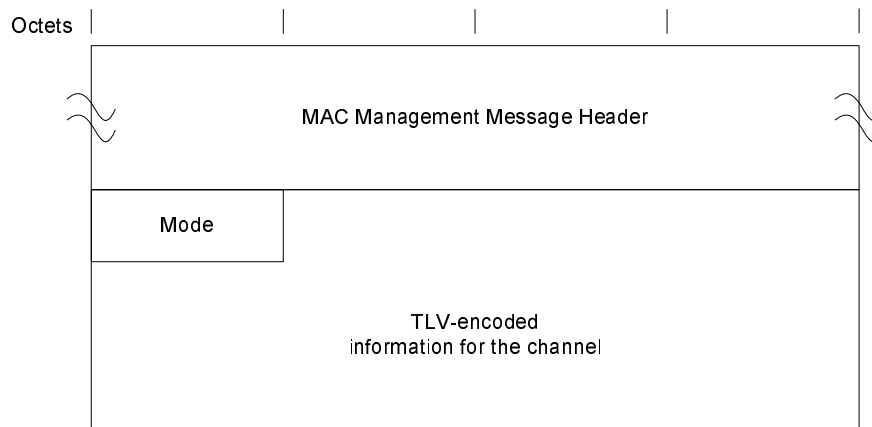


Figure 6-47: Test Request

The parameters of a TST-REQ transmitted by a CMTS MUST be as follows:

Mode:

- 0 = Disable all test modes and reboot.
- 1 = Transmit a continuous (non-burst) upstream signal at the commanded modulation rate, carrier frequency and power level. The chip sequence at the spreader output is replaced with an alternating binary sequence (1, -1, 1, -1, 1, -1, etc.) at nominal amplitude, equal on I and Q. The CM tracks the downstream symbol clock and uses it to generate the upstream symbol clock as in normal synchronous operation.
- 2 = Transmit a continuous (non-burst), unmodulated (CW) upstream signal at the commanded carrier frequency, modulation rate and power level. This is equivalent to replacing the chip sequence at the spreader output with the constant sequence (1, 1, 1, 1, 1, 1, etc.) at nominal amplitude, equal on both I and Q. The CM tracks the downstream symbol clock and uses it to generate the upstream symbol clock as in normal synchronous operation.

In normal operation, the CM MUST ignore any TST-REQ message it receives subsequent to receiving the first SYNC message. Note that this makes it less convenient to use this test mode with a CMTS, since the CMTS may send a SYNC message before the CM sees a TST-REQ message.

After acquiring a downstream signal and prior to receiving a SYNC message, if the CM receives a TST-REQ message (either unicast to the CM itself or broadcast) with Mode != 0, the CM MUST begin the test mode indicated in the Mode parameter, using the channel parameters included in the TST-REQ message.

In test mode, if the CM receives a TST-REQ message (either unicast to the CM itself or broadcast) with Mode = 0, the CM MUST reboot. The CM MUST reboot after the expiration of the T16 timer, a 30 minute test mode timer.

The TST-REQ message MUST be generated by the CMTS in the format shown in figure 6-47, including all of the parameters coded as TLV multiples defined in table 6-31.

Table 6-31: Channel TLV Parameters

Name	Type (1 byte)	Length (1 byte)	Value (variable length)
Modulation Rate	1	1	Multiples of base rate of 160 kHz (Value is 1, 2, 4, 8, 16 or 32). This TLV is present if the test mode is 1.
Frequency	2	4	Upstream carrier frequency (Hz). This TLV is present if the test mode is 1 or 2.
Power	3	1	This TLV specifies the power (unsigned 8-bit, dBmV units) at which the CM transmits. This TLV is present if the test mode is 1 or 2.
S-CDMA US ratio numerator 'M'	4	2	The numerator (M) of the M/N ratio relating the downstream symbol clock to the upstream modulation clock. This TLV is present if the test mode is 1. This TLV is present if the test mode is 2 and the operation is synchronous.
S-CDMA US ratio denominator 'N'	5	2	The denominator (N) of the M/N ratio relating the downstream symbol clock to the upstream modulation clock. This TLV is present if the test mode is 1. This TLV is present if the test mode is 2 and the operation is synchronous.
Upstream channel ID	6	1	The upstream channel ID on which to perform the specified test. This can be included multiple times, once for each channel to be tested. This TLV is present when unicast to a DOCSIS 3.0 CM, if the test mode is 1 or 2. If received by a DOCSIS 3.0 CM as a broadcast message, the CM performs the test on all of the upstream channels in its transmit channel set.

6.4.27 Downstream Channel Descriptor (DCD)

The format and usage of the DCD message is defined in [5].

6.4.28 MAC Domain Descriptor (MDD)

A CMTS MUST transmit an MDD message periodically on every downstream channel in the MAC Domain. The CMTS MUST observe the MDD Interval specified in annex B. The CMTS MUST transmit a separate MDD message for every downstream channel. The CMTS MUST NOT transmit an MDD message with a total Management Message Payload size of more than 8 000 bytes.

The MDD is intended primarily for use by the CM during initialization (see clause 10.2). It also includes parameters related to CM-STATUS reporting which may be useful after registration. During initialization, the CM MUST use the first valid complete MDD (i.e. with all fragments present) received on its selected candidate Primary Downstream Channel as its source for all parameters to be learned from MDD TLVs. All fragments collected need to have the same source MAC address and the same change count. If a CM collects an MDD fragment for the same MAC domain with a change count that is different from that of the fragments already collected, then it MUST discard all previously collected fragments and resume collecting only fragments with the new change count. Also, during initialization, the CM MUST ignore any MDD TLV parameters received in MDD messages on downstream channels other than its selected candidate Primary Downstream Channel.

After registration, the CM MUST use the TLVs applicable to CM-STATUS reporting to control its CM-STATUS reporting as specified in clause 6.4.34. The CM MUST NOT modify anything other than its CM-STATUS reporting behavior in response to changes in the MDD message. For example, the CM does not delete a channel from its Receive Channel Set if that channel is no longer listed in the MDD. The CM MUST ignore any MDD messages received with a source MAC address that is different than the MAC domain address learned during initialization. The CM MUST ignore any changes resulting in a new change count for an MDD message on any of its non-primary channels.

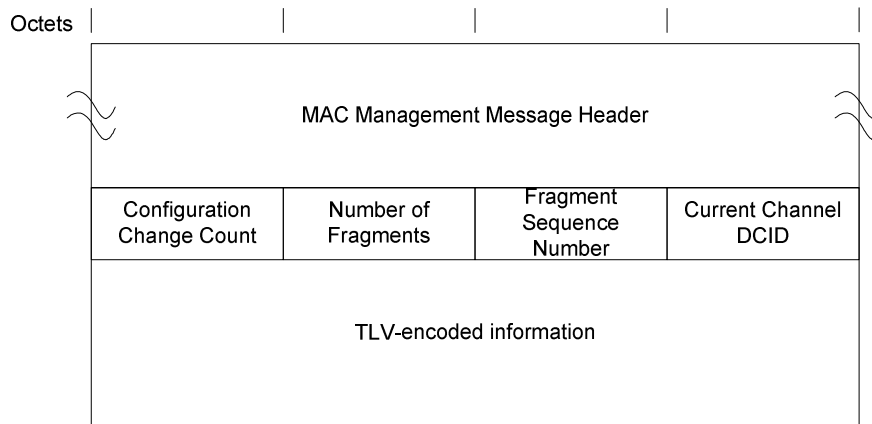


Figure 6-48: MAC Domain Descriptor

A CMTS MUST generate the MDD message in the format shown in figure 6-48, including the following parameters as defined below:

Configuration Change Count: The CMTS increments this field by 1 whenever any of the values in this message change relative to the values in the previous MDD message sent on this downstream channel.

Number of Fragments: Fragmentation allows the MDD TLV parameters to be spread across more than one DOCSIS MAC Frame, thus allowing the total number of MDD TLV parameters to exceed the maximum payload of a single MAC management frame (subject to the constraint stated above). The value of this field represents the number of MDD MAC management frames that a unique and complete set of MDD TLV parameters are spread across to constitute the MDD message. This field is an 8-bit unsigned integer.

Fragment Sequence Number: This field indicates the position of this fragment in the sequence that constitutes the complete MDD message. Fragment Sequence Numbers start with the value of 1 and increase by 1 for each fragment in the sequence. Thus, the first MDD message fragment has a Fragment Sequence Number of 1 and the last MDD message fragment has a Fragment Sequence Number equal to the Number of Fragments. The CMTS MUST NOT fragment any top level TLVs of an MDD. Each MDD message fragment is a complete DOCSIS frame with its own CRC. Other than the Fragment Sequence Number, the framing of one MDD message fragment is independent of the framing of another MDD message fragment. This potentially allows the cable modem to process fragments as they are received rather than reassembling the entire payload. This field is an 8-bit unsigned integer.

Current Channel DCID: The identifier of the downstream channel on which this message is being transmitted.

All other parameters are encoded as TLV tuples, where the type and length fields are each one octet.

6.4.28.1 TLV Encodings

The CMTS MUST use the type values defined in this clause. These are unique within the MDD message, but not across the entire MAC message set. Unless explicitly indicated otherwise, each of these TLVs MUST be included by the CMTS exactly once in each MDD message on a primary-capable downstream channel.

6.4.28.1.1 Downstream Active Channel List TLV

Each instance of this TLV represents one downstream channel in the MAC Domain. The CMTS MAY include this TLV more than once in a given MDD message.

When sending this message on a primary-capable downstream channel, the CMTS MUST include a Downstream Active Channel List TLV for every downstream channel in every MD-DS-SG that contains the current channel.

When sending this message on a non-primary-capable downstream, the CMTS MAY include a Downstream Active Channel List TLV for any primary-capable downstream channel in any MD-DS-SG that contains the current channel. The CMTS SHOULD NOT include a Downstream Active Channel List TLV for non-primary-capable downstreams in a MDD message on a non-primary-capable downstream. The intent is to allow CMs optionally to use the channel list to speed scanning for a primary-capable channel.

The CMTS MUST comply with tables 6-32 and 6-33 for the Downstream Active Channel List TLV.

Table 6-32: Field definitions for Downstream Active Channel List TLV

Type	Length	Value
1	Total number of bytes (including type and length) contained in all sub-TLVs.	Contains sub-TLVs as defined in table 6-33. Each sub-TLV has a one-byte "type" field and one-byte "length" field.

Table 6-33: Sub-TLVs for Downstream Active Channel List TLV

Type	Length	Value
1.1	1	Channel ID: 1 byte. The Downstream Channel ID of the channel being listed.
1.2	4	Frequency: 4 bytes. The center frequency of the downstream channel (Hz). This TLV is intended only to assist CMs in speeding the acquisition of new channels prior to the completion of registration.
1.3	1	Modulation Order/Annex: 1 byte. The CMTS MAY include this TLV. This TLV contains two 4-bit fields: Bits 7 - 4: J.83 annex: 0 = J.83 [i.2] Annex A. 1 = J.83 [i.3] Annex B. 2 = J.83 [i.4] Annex C. 3 to 15 = Reserved. Bits 3 - 0: Modulation Order: 0 = 64-QAM. 1 = 256-QAM. 2 to 15 = Reserved. This TLV is intended only to assist CMs in speeding the acquisition of new channels prior to the completion of registration.
1.4	1	Primary capable: 1 byte. 0 = channel is not primary-capable. 1 = channel is primary-capable. 2 to 255 = Reserved. This TLV is intended only to assist CMs in speeding the acquisition of new channels prior to the completion of registration.
1.5	2	CM-STATUS Event Enable Bitmask: 2 bytes. Each bit in this field represents the enable/disable for a particular event for which status may be reported via the CM-STATUS message. If a bit is 1, CM-STATUS reporting is enabled for the corresponding event. The CMTS MAY include this TLV. If a bit is zero, CM-STATUS reporting is disabled for the corresponding event. If the TLV is omitted then all events are disabled. The details of CM-STATUS message functionality are described in clause 10.4.3. The following bit fields are defined: 0 - Reserved (unused). 1 - MDD timeout. 2 - QAM/FEC lock failure. 3 - Reserved (used for non-channel-specific events). 4 - MDD Recovery. 5 - QAM/FEC Lock Recovery. 6 to 8 - Reserved (used for upstream specific events). 9 to 10 - Reserved (used for non-channel-specific events). 11 to 15 - reserved for future use.
1.6	1	MAP and UCD Transport Indicator: 1 byte. 0 = channel can not carry MAPs and UCDs for the MAC domain for which the MDD is sent. 1 = channel can carry MAPs and UCDs for the MAC domain for which the MDD is sent. 2 to 255 = Reserved. This TLV tells CMs which downstream channels might contain MAPs and UCDs for the MAC domain for which the MDD is sent.

6.4.28.1.2 MAC Domain Downstream Service Group (MD-DS-SG) TLV

When Downstream Channel Bonding is enabled in a MAC Domain, the CMTS MUST transmit one or more instances of this TLV on the primary-capable downstream channels of the MAC Domain. When Downstream Channel Bonding is not enabled in a MAC Domain, the CMTS MAY omit this TLV. When present, CMTS MUST insert this TLV once for each MD-DS-SG reached by this primary-capable downstream channel. The CMTS MUST NOT transmit this TLV on non-primary capable downstream channels. Within each MD-DS-SG encoding, the CMTS SHOULD list only those downstream channels which are relevant to the CM downstream ambiguity process described in clause 10.2.3.

The CMTS MUST comply with tables 6-34 and 6-35 for the MAC Domain Downstream Service Group TLV.

Table 6-34: MAC Domain Downstream Service Group TLV

Type	Length	Value
2	total number of bytes (including type and length) contained in all sub-TLVs.	Contains sub-TLVs as defined in table 6-35. Each sub-TLV has a one-byte "type" field and one-byte "length" field.

Table 6-35: Sub-TLVs for MAC Domain Downstream Service Group TLV

Type	Length	Value
2.1	1	MD-DS-SG identifier (MD_DS_SG_ID): a one-byte value used by the CMTS to identify an MD-DS-SG. For usage details, see clause 10.2.3.
2.2	N (where N = 1 byte for each downstream channel being listed)	Each byte of this field contains a downstream channel ID (DCID) for a different downstream channel which is part of this MD-DS-SG.

6.4.28.1.3 Downstream Ambiguity Resolution Frequency List TLV

This TLV lists downstream frequencies to be used for CM-SG ambiguity resolution per clause 10.2.3. The CMTS MUST include this TLV when sending an MDD message on a primary-capable downstream channel if either Upstream Channel Bonding or Downstream Channel Bonding is enabled for the MAC Domain and this MDD message contains more than one instance of the MD-DS-SG TLV (TLV 2). The CMTS is not required to include this TLV if only one instance of the MD-DS-SG TLV is present.

When this TLV is present, the CMTS MUST list at least one frequency. This TLV indicates to the modem which frequencies it should attempt to receive for downstream service group resolution and in what order. In some topologies, service group resolution efficiency may be improved if the CMTS lists first those frequencies which are most likely to resolve ambiguity. See clause 10.2.3 for details on the service group resolution process. When sending an MDD message on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV.

The CMTS MUST comply with table 6-36 for the Downstream Ambiguity Resolution Frequency List TLV.

Table 6-36: Downstream Ambiguity Resolution Frequency List TLV

Type	Length	Value
3	N (where N = 4 bytes times number of frequencies listed)	Consists of concatenated 4-byte fields. Each 4-byte field contains a center frequency in Hz. The CMTS MUST provide a value which is a multiple of 62,500 Hz. The CM uses these frequencies for downstream CM-SG ambiguity resolution per clause 10.2.3.

6.4.28.1.4 Receive Channel Profile Reporting Control TLV

This TLV controls the reporting of Receive Channel Profiles by CMs in the REG-REQ-MP message. See clause 8.2.4 for details on Receive Channel Profiles. When sending an MDD message on a primary-capable downstream channel, the CMTS MUST include this TLV. The CMTS MUST comply with tables 6-37 and 6-38. When sending an MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV.

Table 6-37: Receive Channel Profile Reporting Control TLV

Type	Length	Value
4	Total number of bytes (including type and length) contained in all sub-TLVs	Contains sub-TLVs as defined in table 6-38. Each sub-TLV has a one-byte "type" field and one-byte "length" field.

Table 6-38: Sub-TLVs for Receive Channel Profile Reporting Control TLV

Type	Length	Value
4.1	1	RCP Center Frequency Spacing. 1 byte: 0 = CM MUST report only Receive Channel Profiles assuming 6 MHz center frequency spacing. 1 = CM MUST report only Receive Channel Profiles assuming 8 MHz center frequency spacing. 2 to 255 = Reserved.
4.2	1	Verbose RCP reporting. 1 byte: 0 = CM MUST NOT provide verbose reporting of all its Receive Channel Profile(s) (both standard profiles and manufacturer's profiles). 1 = CM MUST provide verbose reporting of Receive Channel Profile(s) (both standard profiles and manufacturer's profiles). 2 to 255 = Reserved.
4.3	1	Fragmented RCP transmission. 1 byte: 1 = CM MAY transmit Receive Channel Profile (s) requiring fragmentation (RCPs in excess of 255 bytes) in addition to those that do not. 0,2 to 255 = Reserved. If this sub-TLV is absent from the MDD message, then the CM MUST NOT transmit RCPs requiring fragmentation (see note).
NOTE: At a minimum, CLAB-6M-004 will always be sent for 6 MHz center frequency spacing and CLAB-8M-004 will be sent for 8 MHz center frequency spacing.		

6.4.28.1.5 IP Initialization Parameters TLV

This TLV is used to communicate to the CM certain parameters related to the initialization of the CM's IP-layer services. When sending an MDD message on a primary-capable downstream channel, the CMTS MUST include this TLV. The CMTS MUST comply with tables 6-39 and 6-40. When sending an MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV.

Table 6-39: IP Initialization Parameters TLV

Type	Length	Value
5	Total number of octets (including type and length) contained in all sub-TLVs.	Contains sub-TLVs as defined in table 6-40. Each sub-TLV has a one-octet "type" field and one-octet "length" field.

Table 6-40: Sub-TLVs for IP Initialization Parameters TLV

Type	Length	Value
5.1	1	IP Provisioning Mode (see clause 10.2.5): 0 = IPv4 Only. 1 = IPv6 Only. 2 = Alternate (APM). 3 = Dual-stack (DPM). 4 to 255 = Reserved. The CMTS MUST include this sub-TLV. The CM uses this sub-TLV as defined in clause 10.2.5.
5.2	3	Pre-Registration DSID. Three bytes: bits 23 - 20: Reserved (set to zero). bits 19 - 0: DSID value to be used by the CM for filtering and forwarding Downstream Link-Local Multicast used for IPv6 stack initialization and Neighbor Solicitation prior to registration (see clause 9.2.2). If the CMTS transmits any other IP Initialization Parameter sub-TLVs with a value other than zero and the CMTS enables Multicast DSID Forwarding to any CM on the MAC domain, then the CMTS MUST include this sub-TLV. If the CMTS disables Multicast DSID Forwarding for all CMs in the MAC domain, the CMTS MUST NOT include this sub-TLV.

6.4.28.1.6 Early Authentication and Encryption (EAE) Enable/Disable TLV

This TLV is used to indicate whether the CM is required to perform early authentication and encryption for security purposes. When sending the MDD on a primary-capable downstream channel, the CMTS MUST include this TLV. When sending the MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV.

The CMTS MUST comply with table 6-41. See [15] for additional details.

Table 6-41: Early Authentication and Encryption (EAE) Enable/Disable TLV

Type	Length	Value
6	1	One byte: 0 = early authentication and encryption disabled; 1 = early authentication and encryption enabled; 2 to 255 = Reserved.

6.4.28.1.7 Upstream Active Channel List TLV

Each instance of this TLV represents one active upstream channel in the MAC Domain. The CMTS MAY include this TLV more than once in a given MDD message.

When sending the MDD on a primary-capable downstream channel, the CMTS MUST include an instance of this TLV for every active upstream channel in each MD-CM-SG that includes this downstream channel. When sending the MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV.

The CMTS MUST comply with tables 6-42 and 6-43.

Table 6-42: Field definitions for Active Upstream Channel List TLV

Type	Length	Value
7	Total number of bytes (including type and length) contained in all sub-TLVs.	Contains sub-TLVs as defined in table 6-43. Each sub-TLV has a one-byte "type" field and one-byte "length" field.

Table 6-43: Sub-TLVs for Active Upstream Channel List TLV

Type	Length	Value
7.1	1	The upstream channel ID for a channel being listed.
7.2	2	CM-STATUS Event Enable Bitmask: 2 bytes. Each bit in this field represents the enable/disable for a particular event for which status may be reported via the CM-STATUS message. If a bit is 1, CM-STATUS reporting is enabled for the corresponding event. The CMTS MAY include this TLV. If a bit is zero, CM-STATUS reporting is disabled for the corresponding event. If the TLV is omitted, then all events listed below are disabled. The details of CM-STATUS message functionality are described in clause 10.4.3. The following bit fields are defined: 0 Reserved (unused). 1 to 2 Reserved (used for downstream specific events). 3 Reserved (used for non-channel-specific events). 4 to 5 Reserved (used for downstream specific events). 6 T4 timeout. 7 T3 re-tries exceeded. 8 Successful ranging after T3 re-tries exceeded. 9 to 10 Reserved (used for non-channel-specific events). 11 to 15 Reserved for future use.

6.4.28.1.8 Upstream Ambiguity Resolution Channel List TLV

This TLV lists upstream channel IDs to be used for CM-SG ambiguity resolution per clause 10.2.3. When sending the MDD on a primary-capable downstream channel, the CMTS MUST include this TLV. When sending the MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV. The CMTS MUST comply with table 6-44. The CMTS MUST list at least one channel ID in the Upstream Ambiguity Resolution Channel List for each MD-US-SG served by that MDD message. The CM will choose a channel from this list for its initial ranging attempt per clause 10.2.3.

Table 6-44: Upstream Ambiguity Resolution Channel List TLV

Type	Length	Value
8	N (where N = the number of channel IDs listed).	Each byte of this field contains an upstream channel ID (UCID) for a channel being listed.

6.4.28.1.9 Upstream Frequency Range TLV

This TLV indicates the frequency range of the plant reserved for upstream transmission. When sending the MDD on a primary-capable downstream channel, the CMTS MUST include this TLV. When sending the MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV. The CMTS MUST format and use the TLV as indicated in table 6-45.

The CM MUST configure its upstream transmitters prior to initial ranging to use the highest Upstream Frequency Range setting that it supports which is not greater than the capability advertised by the CMTS.

Table 6-45: Upstream Frequency Range TLV

Type	Length	Value
9	1	Upstream Frequency Range: 1 byte: 0 = Standard Upstream Frequency Range (see [12]). 1 = Extended Upstream Frequency Range (see [12]). 2 to 255 = Reserved.

6.4.28.1.10 Symbol Clock Locking Indicator

EN 302 878-3 [4] requires the CMTS to lock its Symbol Clock to the Master Clock. This TLV indicates whether or not the symbol clock for the current downstream channel is locked to the CMTS Master Clock. When sending the MDD on a primary-capable downstream channel, the CMTS MUST include this TLV. When sending the MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV. The CMTS MUST comply with table 6-46. If this TLV is not present, the MDD MUST be considered invalid by the CM.

Table 6-46: Symbol Clock Locking Indicator TLV

Type	Length	Value
10	1	Symbol Clock Locking Indicator: 0 = Symbol Clock is not locked to Master Clock. 1 = Symbol Clock is locked to Master Clock.

6.4.28.1.11 CM-STATUS Event Control

The CM-STATUS reporting mechanism includes a random holdoff prior to transmission of status report messages. This TLV indicates the value of that random holdoff timer to be used by the CM when determining when/whether to transmit a CM-STATUS message. This TLV associates a separate hold-off timer value with each CM-STATUS event type code managed by the CMTS. When the CM receives an MDD message on its Primary Downstream Channel that does not include an Event Control Encoding for an event type, the CM does not transmit CM-STATUS messages with that event type code. A valid MDD message may have any number of CM-STATUS Event Control Encodings as long as each event code is unique.

Event reporting is enabled jointly by the presence of the appropriate Event Control TLV and the appropriate bit in the CM-STATUS Event Enable Bit Mask TLV 1.5, 7.2 or 15. Refer to clause 10.4.3 for requirements on enabling event reportings.

The CMTS MAY include one instance of this TLV in a MDD message on a primary-capable downstream channel. When sending an MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV. The CMTS MUST comply with table 6-47.

Table 6-47: CM-STATUS Event Control TLV

Type	Length	Value
11	10	Event Control Encoding. A valid encoding contains a single instance of each of the subtypes defined below.
11.1	1	Event Type Code as defined in table 10-3.
11.2	2	Maximum Event Holdoff Timer in units of 20 milliseconds. Valid range: 1..65535.
11.3	1	Maximum Number of Reports per event: 0: Unlimited number of reports. 1 to 255: Maximum number of reports for an event type reporting transaction.

The CM MUST silently ignore event type codes unknown to the CM. The CM MUST silently ignore unknown subtypes of an Event Control Encoding and implement its known subtypes.

6.4.28.1.12 Upstream Transmit Power Reporting

This TLV indicates whether the CM should report its upstream transmit power in the SSAP field of the MAC Management Header of the RNG-REQ, INIT-RNG-REQ and B-INIT-RNG-REQ messages. The reporting of upstream transmit power is described in clause 6.4.5. When sending the MDD on a primary-capable downstream channel, the CMTS MAY include this TLV. If the CMTS does not include this TLV with a value indicating transmit power reporting enabled, it MUST NOT provision any CM with a Transmit Channel Set containing more than one channel. When present, this TLV MUST be formatted as shown in table 6-48. When sending the MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV.

Table 6-48: Upstream Transmit Power Reporting TLV

Type	Length	Value
12	1	0: CM does not report transmit power in RNG-REQ, INIT-RNG-REQ and B-INIT-RNG-REQ messages. 1: CM reports transmit power in RNG-REQ, INIT-RNG-REQ and B-INIT-RNG-REQ messages. 2 - 255: Reserved.

6.4.28.1.13 DSG DA-to-DSID Association Entry

This TLV conveys the association between a DSID and a MAC Destination Address being used for DSG. It is necessary to communicate this information in a broadcast downstream message for DOCSIS 3.0 DSG modems operating in one-way mode. The CMTS is not required to include this TLV in the MDD if the CMTS has been configured to disable Multicast DSID Forwarding on a Global or Mac Domain basis. When sending the MDD on a primary-capable downstream channel, the CMTS includes this TLV if DCD messages are also being sent on the downstream channel. The CMTS includes one instance of this TLV for each multicast MAC DA in the DCD message. The CMTS may include one instance of this TLV for each unicast MAC DA in the DCD message. The CMTS does not use a given DSID value in more than one instance of this TLV. When sending the MDD on a non-primary-capable downstream channel, the CMTS does not include this TLV. The format and contents of this TLV are detailed in tables 6-49 and 6-50.

Table 6-49: DSG DA-to-DSID Association Entry TLV

Type	Length	Value
13	Total number of bytes (including type and length) contained in all sub-TLVs.	Contains sub-TLVs as defined in table 6-50. Each sub-TLV has a one-byte "type" field and one-byte "length" field. Each sub-TLV appears exactly once.

Table 6-50: Sub-TLVs for DSG DA-to-DSID Association Entry TLV

Type	Length	Value
13.1	6	DA: the 48-bit MAC DA to which this association applies.
13.2	3	Bits 23-20: Reserved. Bits 19-0: the 20-bit DSID associated with the DA contained in sub-TLV 13.1.

6.4.28.1.14 CM-STATUS Event Enable for Non-Channel-Specific Events

The CMTS MAY include one instance of this TLV in a MDD message on a primary-capable downstream channel. When sending an MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV.

Table 6-51: CM-STATUS Event Enable for Non-Channel-Specific Events TLV

Type	Length	Value
15	2	<p>CM-STATUS Event Enable Bitmask for Non-Channel-Specific Events; 2 bytes.</p> <p>Each bit in this field represents the enable/disable for a particular non-channel-specific event for which status may be reported via the CM-STATUS message. If a bit is 1, CM-STATUS reporting is enabled for the corresponding event. If a bit is zero, CM-STATUS reporting is disabled for the corresponding event. If the TLV is omitted, then all events listed below are disabled. The details of CM-STATUS message functionality are described in clause 10.4.3.</p> <p>The following bits are defined:</p> <ul style="list-style-type: none"> 0 Reserved (unused). 1 to 2 Reserved (used for downstream specific events). 3 Sequence out-of-range. 4 to 5 Reserved (used for downstream specific events). 6 to 8 Reserved (used for upstream specific events). 9 CM operating on battery backup. 10 CM returned to A/C power. 11 to 15 Reserved for future use.

6.4.28.1.15 Extended Upstream Transmit Power Support

This encoding within the MDD message signals whether or not modems may transmit at power levels greater than the default P_{\max} values defined in [12] prior to registration (post registration behavior is controlled via the Extended Upstream Transmit Power capability as defined in clause C.1.3.1.40). By default, the CMTS MUST set this TLV to On unless a mechanism is provided to administratively configure this setting on and off. When this TLV is present and set to On, the CM is permitted to exceed the default P_{\max} values as specified in [12] prior to registration.

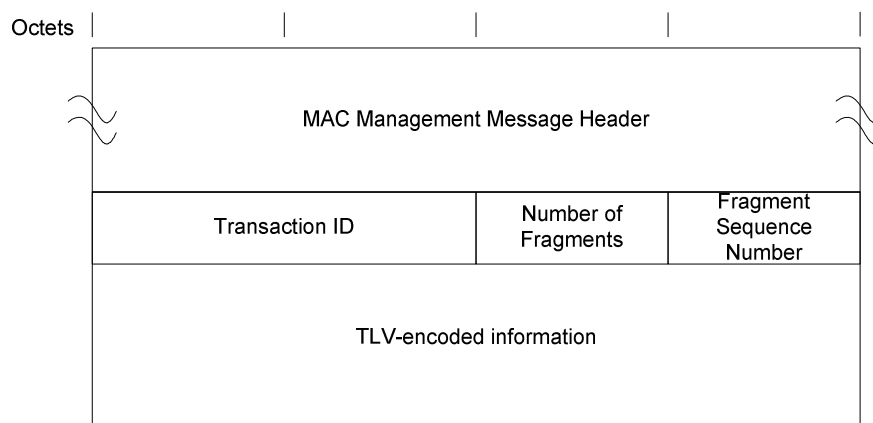
The CMTS MUST include one instance of this TLV in an MDD message on a primary-capable downstream channel. When sending an MDD on a non-primary-capable downstream channel, the CMTS MUST NOT include this TLV.

Table 6-52: Extended Upstream Transmit Power Support

Type	Length	Value
16	1	<p>Extended Upstream Transmit Power Support: 1 byte:</p> <ul style="list-style-type: none"> 0 Extended Upstream Transmit Power Support Off. 1 Extended Upstream Transmit Power Support On. 2 to 255 Reserved.

6.4.29 Dynamic Bonding Change Request (DBC-REQ)

A Dynamic Bonding Change Request message is transmitted by the CMTS in order to change upstream and/or downstream bonding parameters or downstream multicast parameters. Only one DBC transaction per CM can be in the process at any time. The CMTS MUST wait for any ongoing transaction for a particular CM to be finished before a new transaction can be initiated with that CM. The DBC-REQ message is formatted as shown in figure 6-49.

**Figure 6-49: Dynamic Bonding Change Request Message**

The Parameters for a DBC-REQ transmitted by a CMTS MUST be as follows:

Transaction ID: Unique identifier for this transaction assigned by the CMTS.

Number of Fragments: Fragmentation allows the DBC-REQ TLV parameters to be spread across more than one DOCSIS MAC Frame, thus allowing the total number of DBC-REQ TLV parameters to exceed the maximum payload of a single MAC management frame. The value of this field represents the number of DBC-REQ MAC management frames that a unique and complete set of DBC-REQ TLV parameters are spread across to constitute the DBC-REQ message. This field is an 8-bit unsigned integer. The default value for this field is 1.

Fragment Sequence Number: This field indicates the position of this fragment in the sequence that constitutes the complete DBC-REQ message. Fragment Sequence Numbers start with the value of 1 and increase by 1 for each fragment in the sequence. Thus, the first DBC-REQ message fragment would have a Fragment Sequence Number of 1 and the last DBC-REQ message fragment would have a Fragment Sequence Number equal to the Number of Fragments. The CM is not required to reorder DBC message fragments. The CMTS MUST ensure that the message fragments arrive in order at the CM either by sending all message fragments on a single downstream or by transmitting fragments such that individual channel latencies do not affect fragment order. The CMTS MUST NOT fragment within any top level TLVs. Each DBC-REQ message fragment is a complete DOCSIS frame with its own CRC. Other than the Fragment Sequence Number, the framing of one DBC-REQ message fragment is independent of the framing of another DBC-REQ message fragment. This field is an 8-bit unsigned integer. The default value for this field is 1.

All other parameters are coded as TLV tuples as defined in annex C. A DBC-REQ transmitted by a CMTS MUST contain at least one of the following:

Transmit Channel Configuration: Specification of the rules to be used to make changes to a Transmit Channel Set (refer to clause C.1.5.1).

Service Flow SID Cluster Assignments: Specification of the rules to be used to make changes to a Service Flow Cluster Assignments (refer to clause C.1.5.2).

Receive Channel Configuration: Specification of the rules to be used to make changes to a Receive Channel Set (refer to clause C.1.5.3).

DSID Encodings: Specification of the rules to be used to make changes to a DSID (refer to clause C.1.5.3.8).

Security Association Encodings: Specification of the rules to be used to make changes to a SAID (refer to clause C.1.5.5).

If Privacy is enabled, the CMTS MUST also format the DBC-REQ message to contain:

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the DBC-REQ message's Attribute list (refer to clause C.1.4.1). In the case of a fragmented DBC-REQ message, the HMAC-Digest appears only once as the final Attribute in the last fragment of the DBC-REQ message.

6.4.30 Dynamic Bonding Change Response (DBC-RSP)

The CM MUST transmit a Dynamic Bonding Change Response in response to a received Dynamic Bonding Change Request (DBC-REQ) message. The DBC-RSP message transmitted by a CM MUST be formatted as shown in figure 6-50.

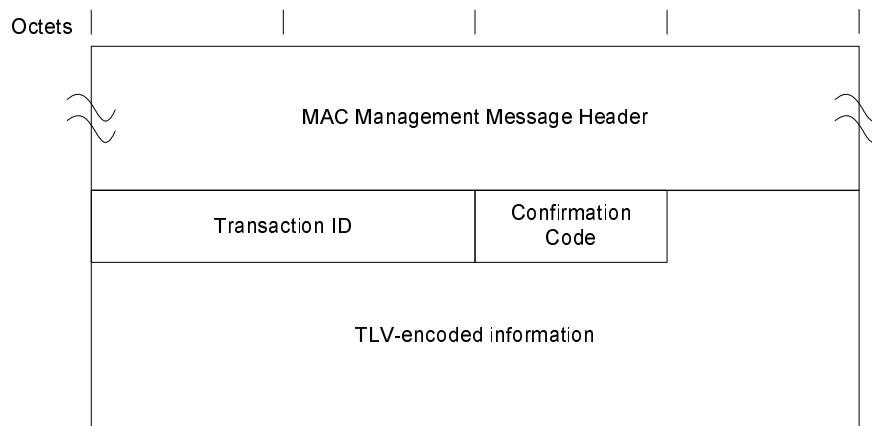


Figure 6-50: Dynamic Bonding Change Response Message

The Parameters of the DBC-RSP transmitted by a CM MUST be as follows:

Transaction ID: Transaction ID from the corresponding DBC-REQ.

Confirmation Code: An 8-bit Confirmation Code. The valid codes are defined in clause C.4.

All other parameters are encoded as TLV tuples as defined in annex C.

If the transaction is unsuccessful due to TCC Encodings or RCC Encodings and the Confirmation Code is not one of the major error codes in annex C, the DBC-RSP transmitted by the CM MUST contain at least one of the following as defined below:

TCC Error Set: A TCC Error Set and identifying TCC Reference is included for at least one failed TCC in the corresponding DBC-REQ. Every TCC Error Set includes at least one specific failed parameter of the corresponding TCC. It does not need to include every failed parameter of the corresponding TCC. This parameter is omitted if the entire DBC-REQ is successful (refer to clause C.1.5.1).

RCC Error Set: An RCC Error Set. This parameter is included to report an error in an RCC encoding in the corresponding DBC-REQ. Every RCC Error Set includes at least one specific failed parameter of the corresponding RCC. It does not need to include every failed parameter of the corresponding RCC. This parameter is omitted if the entire DBC-REQ is successful (refer to clause C.1.5.3).

In the case where the CM is unable to acquire one or more of the upstream and/or downstream channels assigned via the TCC and/or RCC encodings (respectively), the CM needs to report back to the CMTS the list of channels that it was unable to acquire so that the CMTS can take appropriate action. If the CM is unable to acquire one or more of the downstream channels assigned to it in the RCC, the CM MUST include an RCC encoding with a Partial Service Downstream Channels TLV in the DBC-RSP, which includes a list of the downstream channels that could not be acquired. If the CM is unable to acquire one or more of the upstream channels assigned to it in the TCC, the CM MUST include a TCC encoding with a TCC Error Encoding for each upstream channel it was unable to acquire in the DBC-RSP, corresponding to the TCC encoding that assigned that upstream channel in the DBC-REQ. This is because each TCC encoding describes the actions to take for a single upstream channel. Note that this is different from the case of reporting an error in the encoding, where only a single error needs to be reported (even if multiple errors exist).

Regardless of success or failure, if Privacy is enabled for the CM, the DBC-RSP message transmitted by the CM MUST contain:

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the DBC-RSP message's Attribute list (refer to clause C.1.4.1).

6.4.31 Dynamic Bonding Change Acknowledge (DBC-ACK)

The Dynamic Bonding Change Acknowledge MUST be transmitted by a CMTS in response to a received Dynamic Bonding Change Response (DBC-RSP) message from a CM. The DBC-ACK message is formatted as shown in figure 6-51.

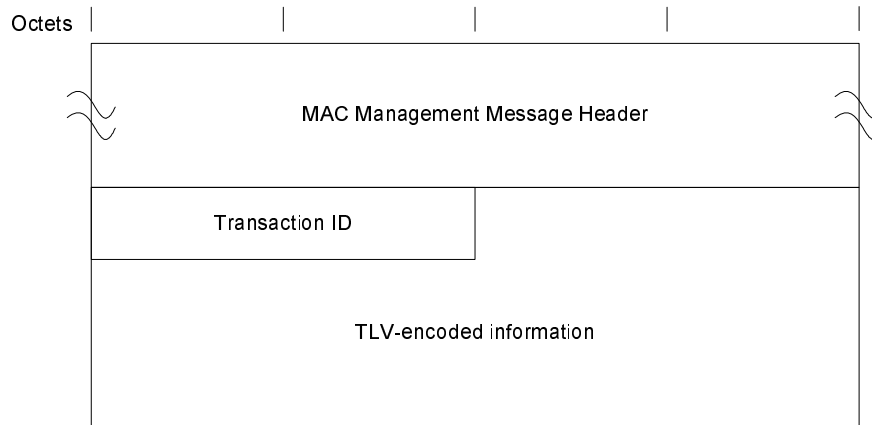


Figure 6-51: Dynamic Bonding Change Acknowledge Message

The parameters of a DBC-ACK message transmitted by a CMTS MUST be as follows:

Transaction ID: Transaction ID from the corresponding DBC-REQ.

If Privacy is enabled, the DBC-ACK message transmitted by the CMTS MUST contain:

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

HMAC-Digest: The HMAC-Digest Attribute is a keyed message digest (to authenticate the sender). The HMAC-Digest Attribute is the final Attribute in the DBC-ACK message's Attribute list (refer to clause C.1.4.1).

6.4.32 DOCSIS Path Verify Request (DPV-REQ)

The DOCSIS Path Verify (DPV) MAC Management Messages are used for measuring latency within the DOCSIS system. This message may be sent to either the DOCSIS multicast MAC address (refer to annex A) or directly to a unicast MAC address of a CM.

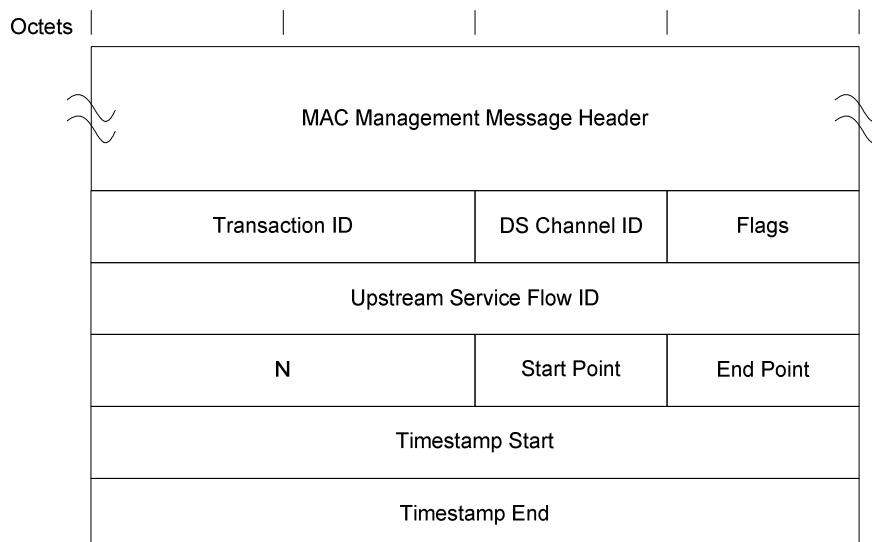


Figure 6-52: DPV-REQ MAC Message

When transmitting a DPV-REQ message, the CMTS MUST use the format shown in figure 6-52, including the following parameters as defined below:

Transaction ID: Unique identifier for this transaction assigned by the CMTS.

DS Channel ID: This is the Channel ID of the DOCSIS downstream channel on which the sender has requested that the measurement take place. It is used to select a DPV Counter Group. If the value of the DC field is non-zero, the CMTS sets this field to indicate a Channel ID in the CM's Receive Channel Set.

Flags: Formatted as follows:

Bits 7 to 6: DC: DPV Statistical Group Control 00 = Do nothing. 01 = Merge latency measurement into Statistical Group #1. 10 = Merge latency measurement into Statistical Group #2. 11 = Clear Statistical Groups #1 and #2.	Bits 5 to 1: Reserved bits. The CMTS sets these bits to 0. The CM MUST ignore these bits.	Bit 0: E: Echo bit. If E=1, the CM MUST send a DPV-RSP message. If E=0, the CM MUST NOT send a DPV-RSP.
---	--	---

US SFID: Upstream Service Flow ID: This is the upstream Service Flow on which the CM should send the DPV-RSP message. If this field is all zeros and the E bit is asserted, then the CM SHOULD use its primary upstream Service Flow.

N: Measurement averaging factor. This value is used by the CM to calculate a running average as described in clause 10.5. If the value of DC is either 01 or 10, the CMTS MUST set this field to a non-zero value.

Start Reference Point: This is the DPV Reference Point from which the DPV measurement originates.

End Reference Point: This is the DPV Reference Point at which the DPV measurement terminates.

Timestamp Start: If the CMTS owns the Start Reference Point, it will place a copy of its local DOCSIS timestamp in this field. Otherwise, the CMTS sets this field to all zeros.

Timestamp End: This value is initialized to all zeros by the CMTS.

The multicast version of the DPV-REQ message is useful when all CMs are passively logging latency measurements without sending a DPV-RSP (E bit not asserted). A multicast message also ensures that all CMs receive the same messages so that CMTS to CM latencies can be more accurately compared.

The CMTS should be cautious about asserting the E bit when sending a multicast DPV-REQ as this will cause all CMs to simultaneously attempt to send a DPV-RSP. This may be a useful technique for measuring upstream access latency during congestion, but there will be an impact to the operational capability of the upstream. The CMTS can use a 3-byte Downstream Service Extended Header (see clause 6.2.5.6) to limit the number of CMs that would receive and potentially respond to a multicast DPV-REQ.

The CMTS MAY support the generation of the DPV-REQ message in the downstream direction. The CM MUST support the reception of the DPV-REQ message in the downstream direction.

6.4.33 DOCSIS Path Verify Response (DPV-RSP)

The DPV MAC Management Messages are used for estimating latency and skew within the DOCSIS system. The CM **MUST** comply with figure 6-53 for DPV Response messages. This message is sent by the CM to the unicast MAC address of the CMTS.

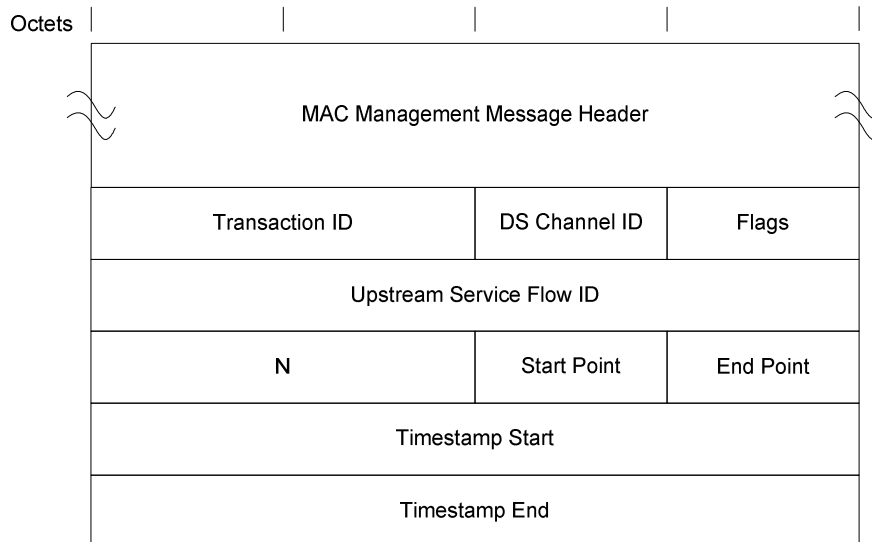


Figure 6-53: DPV-RSP MAC Message

The CM **MUST** copy the values of all fields from the DPV-REQ into the identical fields in the DPV-RSP message with the exception of the following cases:

Timestamp Start: If the CM owns the Start Reference Point, it **MUST** place a copy of its local DOCSIS timestamp in this field. Otherwise, this value is copied from the identical field in the DPV-REQ message.

Timestamp End: If the CM owns the End Reference Point, it **MUST** place a copy of its local DOCSIS timestamp in this field. Otherwise, this value is copied from the identical field in the DPV-REQ message.

The CM **MUST** support the generation of the DPV-RSP message in the upstream direction. The CMTS **MAY** support the reception of the DPV-RSP message in the upstream direction.

6.4.34 Status Report (CM-STATUS)

A CM **MUST** generate the CM-STATUS message compliant with figure 6-54, including the Transaction ID and Event Type. The inclusion of these parameters in the beginning of the message body allows the CMTS to quickly filter events without parsing through the TLV structure.

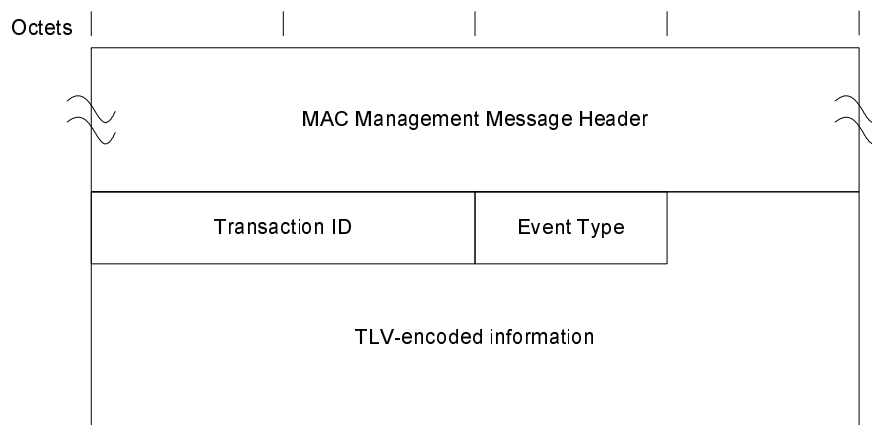


Figure 6-54: CM-STATUS Report

Transaction ID: This is a 2-byte value that identifies a reported transition of the event from off to on. Upon MAC Initialization, the CM MUST report the first CM-STATUS Transaction ID for each event type as 1. The CM MUST NOT use a Transaction ID value of 0 (zero). This ensures that the CMTS can always reset its last received Transaction ID to 0 and be assured of processing the next CM-STATUS message. When incrementing a value of 65 535, the CM wraps around to a value of 1.

Event Type Code: This field contains a unique code which describes the event condition. Refer to table 6-54. The CM MUST include this field.

6.4.34.1 TLV Encodings

Table 6-53: CM-STATUS TLV Encodings

Type	Length	Value
1	N	Status Event. This TLV is repeated for each error event that is being reported by the CM.
1.2	1-80	Event Description. This is an optional vendor specific text string containing details on the failure. The CM MAY include this TLV.
1.4	1	Downstream Channel ID. This is the channel on which the error was detected. It is the same channel ID advertised for the failed channel in MDD messages. The CM-STATUS message includes one instance of this encoding for each channel for which the event type is considered to be "on". This TLV is included for certain status events as indicated in table 10-3.
1.5	1	Upstream Channel ID. This is the channel on which the error was detected. The CM-STATUS message includes one instance of this encoding for each channel for which the event type is considered to be "on". This TLV is included for certain status events as indicated in table 10-3.
1.6	3	DSID. This is the value of the DSID on which the error occurred. The CM-STATUS message includes one instance of this encoding for each DSID for which the event type is considered to be "on". This TLV is included for certain status events as indicated in table 10-3.

6.4.35 CM Control Request (CM-CTRL-REQ)

The CM-CTRL-REQ command is used to enforce specific CM actions. It is a replacement to the DOCSIS 2.0 UP-DIS management message. The CMTS MUST support the CM-CTRL-REQ message. The CM MUST support the CM-CTRL-REQ message.

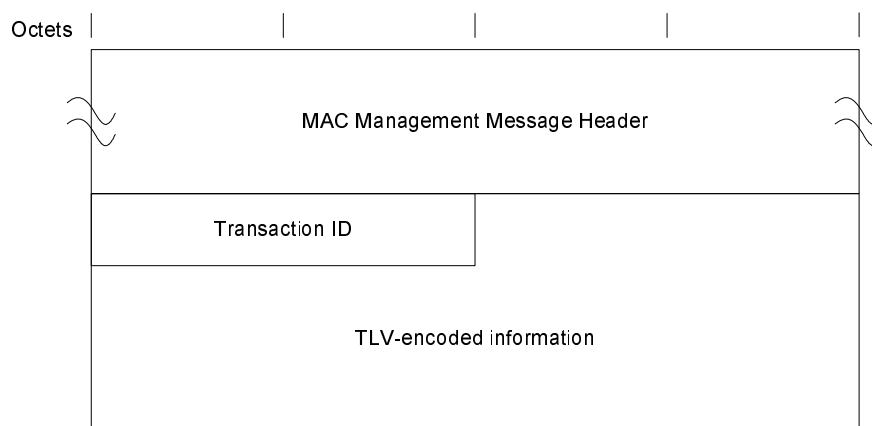


Figure 6-55: CM-CTRL-REQ

A CMTS MUST generate the CM-CTRL-REQ message compliant with figure 6-55 including the following parameter:

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

The CM MUST accept a CM-CTRL-REQ message on any available downstream.

6.4.35.1 TLV Encodings

The CMTS MUST use the TLV encodings described in table 6-54. The CM MUST support each action defined by the TLV encodings described in table 6-54. The CM MUST NOT act upon unknown TLVs in a CM-CTRL-REQ message.

Table 6-54: CM-CTRL-REQ TLV Encodings

Type	Length	Value
1	1	Upstream Channel RF Mute. This field contains the Channel ID of the upstream to mute or un-mute. A value of 0 will mute or un-mute all channels. The mute operation is a low level disabling of the physical layer transmitter that is currently using the channel ID. It will not directly change the MAC layer state, although if the mute period is long enough the MAC layer will experience T4 timeout as if the channel has become physically unavailable. If all channels are muted and the CM encounters a condition which leads it to the Re-Init MAC state, the CM MUST defer re-initialization and remain muted until the mute timer expires, an un-mute command is received or a Lost SYNC event occurs, at which point it performs a re-init MAC and is no longer muted.
2	4	RF Mute Timeout Interval. For the RF Mute operation, this field controls the length of time that the upstream channel(s) are muted. This field is a 32-bit unsigned integer in units of milliseconds. The CMTS MUST include the RF Mute Timeout Interval TLV when the Upstream Channel RF Mute TLV is included in the CM-CTRL-REQ message. A timeout of 0x00000000 is an indication to un-mute the channel(s) immediately. A timeout of 0xFFFFFFFF is an indication to mute the channel(s) indefinitely.
3	1	CM Reinitialize. A value of 1 instructs the CM to reinitialize its MAC with a CM Initialization Reason of CM_CTRL_INIT and will begin a new registration process. Any value other than 1 is ignored.
4	1	Disable Forwarding. A value of 1 will disable forwarding of data PDUs in both the upstream and downstream direction. A value of 0 will enable forwarding of data PDUs in both the upstream and downstream direction. Any value other than 0 or 1 will be ignored.
5	7	Override for the Downstream Status Event Enable Bitmask.
5.1	1	Downstream Channel ID.
5.2	2	Downstream Status Event Enable Bitmask (clause 6.4.28).
6	7	Override for the Upstream Status Event Enable Bitmask.
6.1	1	Upstream Channel ID.
6.2	2	Upstream Status Event Enable Bitmask (clause 6.4.28).
7	2	Override for the CM-STATUS Event Enable Bitmask for Non-Channel-Specific Events (clause 6.4.28).

The CM uses the CM-CTRL-REQ to enforce specific CM actions according to the requirements specified in clause 10.4.3.

6.4.36 CM Control Response (CM-CTRL-RSP)

The CM-CTRL-RSP message is used to confirm receipt of a CM-CTRL-REQ message. The CM MUST send a CM-CTRL-RSP message every time it receives a CM-CTRL-REQ message prior to performing the action described in the CM-CTRL-REQ message.

The CMTS SHOULD consider a previously transmitted CM-CTRL-REQ message to be lost if the CMTS has not received a CM-CTRL-RSP message from the CM within 5 s.

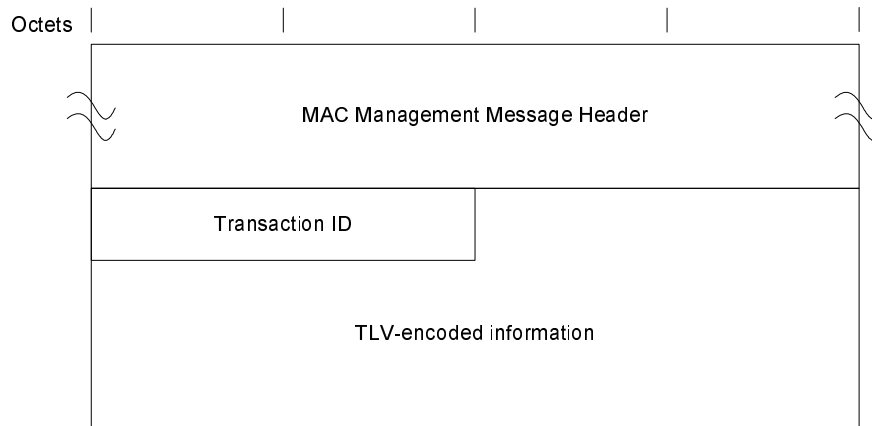


Figure 6-56: CM-CTRL-RSP

A CM MUST generate the CM-CTRL-RSP message compliant with figure 6-56 including the following parameter:

Transaction ID: A 16-bit unique identifier for this transaction from the corresponding CM-CTRL-REQ message.

The TLVs in the CM-CTRL-RSP are the same top-level TLVs that are used in the CM-CTRL-REQ message, except that they are all of length 1 and can only have values of 0 or 1. The CM MUST include every top-level TLV from the CM-CTRL-REQ message in the CM-CTRL-RSP. Each TLV included by the CM in the CM-CTRL-RSP MUST have a length of 1 and either a value of 0 if the CM will apply the TLV (success) or a value of 1 if the CM cannot apply the TLV (fail). The CM MUST include unknown TLVs from the CM-CTRL-REQ message in the CM-CTRL-RSP using a value of 1 (fail).

7 Media Access Control Protocol Operation

7.1 Timing and Synchronization

One of the major challenges in designing a MAC protocol for a cable network is compensating for the delays involved. These delays can be an order of magnitude larger than the transmission burst time in the upstream. To compensate for these delays, the cable modem needs to 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 by the CMTS during a ranging process, for each cable modem.

7.1.1 Global Timing Reference

For TDMA channels, 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.

For S-CDMA channels, the CMTS also creates a global timing reference by transmitting the Time Synchronization (SYNC) and Upstream Channel Descriptor (UCD) MAC messages downstream at a nominal frequency. See [12].

The Transmission Convergence sublayer needs to operate closely with the MAC sublayer to provide an accurate timestamp for the SYNC message.

It is intended that the nominal interval between SYNC messages be tens of milliseconds and the nominal interval between UCD messages be no more than 2 seconds. This imposes relatively little downstream overhead while letting cable modems acquire their global timing synchronization quickly.

For DOCSIS 3.0 CMs, the CMTS conveys the global timing reference to a CM on the CM's Primary Downstream Channel. The CM MUST use a single synchronization timebase obtained on its Primary Downstream Channel for upstream burst timing for all of the upstream channels that the CM is using. A cable modem MUST NOT use an upstream channel until it has successfully synchronized to its Primary Downstream Channel as defined in clause 10.2.1.

7.1.2 CM Synchronization

When the MDD Symbol Clock Locking Indicator TLV is set to "Not Locked", 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 (as defined in [12]). When the MDD Symbol Clock Locking Indicator TLV is set to "Not Locked", the cable modem MUST NOT lock to the downstream symbol clock on its Primary Downstream Channel. In this case, the cable modem acquires the synchronization timebase to be used for upstream burst timing from the SYNC messages.

When the MDD Symbol Clock Locking Indicator indicates "Locked", the cable modem achieves MAC synchronization once it has received at least two SYNC messages, received one UCD message, has locked to the downstream symbol clock and has verified that its clock tolerances are within specified limits (as defined in [12]). When the MDD Symbol Clock Locking Indicator indicates "Locked", the cable modem MUST lock to the downstream symbol clock on its Primary Downstream Channel using the M and N integer frequency ratio values specified in [4] as the source for upstream burst timing, regardless of whether its upstream channels are using TDMA, S-CDMA or both.

When the MDD is not present and the CM selects a TDMA upstream channel, 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 (as defined in [12]). When the MDD is not present and the CM selects a TDMA upstream channel, the cable modem MUST not lock to the downstream symbol clock on its Primary Downstream Channel. In this case, the cable modem acquires the synchronization timebase to be used for upstream burst timing from the SYNC messages.

When the MDD is not present and the CM selects an S-CDMA upstream channel, the cable modem achieves MAC synchronization once it has received at least two SYNC messages, received one UCD message, has locked to the downstream symbol clock and has verified that its clock tolerances are within specified limits (as defined in [12]). When the MDD is not present and the CM selects an S-CDMA upstream channel, the CM MUST lock to the downstream symbol clock on its Primary Downstream Channel using the M and N integer frequency ratio values provided in the UCD as the source for upstream burst timing.

When an MDD is not present and the modem first achieves MAC synchronization on a TDMA channel, in event that the channel is changed to S-CDMA channel, the modem will need to reacquire synchronization and re-range.

7.1.3 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 of the CM and CMTS MUST be relatively constant with the exception of the timing offsets specified in [12], related to modulation rate changes to accommodate a Pre-3.0 DOCSIS upstream receiver implementation. For TDMA, any variation in the PHY delays MUST be accounted for by the CMTS in the guard time of the upstream PMD overhead.

7.1.3.1 Broadcast Initial Ranging

First, a cable modem MUST synchronize to the downstream as described in clause 7.1.2 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 a Broadcast Initial Maintenance Region (refer to clause 7.2.1.2.3). The CMTS MUST schedule Broadcast Initial Maintenance regions large enough to account for the worst case round-trip plant delay. On S-CDMA channels, the CMTS MUST schedule Broadcast Initial Maintenance transmit opportunities such that they align with S-CDMA frames and span an integral number of S-CDMA frames (refer to [12]).

The cable modem MUST transmit either a Bonded Initial Ranging Request message (B-INIT-RNG-REQ), an Initial Ranging Request message (INIT-RNG-REQ) or a Ranging Request message (RNG-REQ) in a Broadcast Initial Maintenance region. The CM MUST transmit a B-INIT-RNG-REQ if the CM detected an MDD on its candidate Primary Downstream Channel and is ranging for the first time after power-up or reinitialization on the first upstream channel (see clause 10.2.3). A CM MUST transmit an INIT-RNG-REQ if the upstream is a Type 3 or a Type 4 channel (which can be determined from the UCD) and the CM is not ranging for the first time following initialization after power-up or reinitialization on the first upstream channel. If none of the conditions for transmitting a B-INIT-RNG-REQ or an INIT-RNG-REQ are met, the CM MUST transmit a RNG-REQ. The CM sets the SID field in the RNG-REQ or INIT-RNG-REQ as defined in clause 6.4.5. The CM MUST set its initial timing offset to the amount of internal fixed delay equivalent to putting this CM next to the CMTS (i.e. no plant delay). This amount includes delays introduced through a particular implementation and the downstream PHY interleaving latency.

Once the CMTS has successfully received the RNG-REQ, INIT-RNG-REQ or B-INIT-RNG-REQ 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 (unless the CM has retained a previous Primary SID during a UCC, DCC or UCD change or a Ranging SID through registration or DBC messaging) until it has completed the registration process. The message from the CMTS MUST also contain information on RF power level adjustment and offset frequency adjustment as well as any timing offset corrections. Ranging adjusts each CM's timing offset such that it appears to be located right next to the CMTS.

7.1.3.2 Unicast Initial Ranging

The cable modem MUST now wait for an individual Station Maintenance or Unicast Initial Maintenance region assigned to its temporary SID (or previous primary SID if ranging as a result of a UCC, DCC or UCD change or Ranging SID if one has been assigned). The CM MUST now transmit a Ranging Request (RNG-REQ) message at this time using the temporary SID (or primary/Ranging SID, as appropriate) 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 by the CM and CMTS, 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 clause 10, for complete details on the entire initialization sequence. In particular, state machines, the applicability of retry counts and timer values for the ranging process are defined in clause 10.3.

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

7.1.4 Timing Units and Relationships

The SYNC message conveys a time reference with a resolution of 6,25/64 microseconds (10,24 MHz) to allow the CM to track the CMTS clock with a small phase offset. Since this timing reference 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 time needed for CM transmission of a fixed number of symbols. 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.

7.1.4.1 TDMA Timing Units and Relationships

7.1.4.1.1 Mini-Slot Capacity

On TDMA channels, 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 7-1 relates mini-slots to the SYNC time ticks (assuming QPSK modulation).

Table 7-1: 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: The symbols/byte is a characteristic of an individual burst transmission, not of the channel. A mini-slot in this instance could represent a minimum of 16 or a maximum of 48 bytes, depending on the modulation choice.

In an upstream channel is a Type 3a or 4a channel, the Mini-slot Size field (M) of the UCD MAY be assigned the value 0 by the CMTS for a 5,12 Msps channel, in which case the mini-slot size is 1 Timebase Tick. If a channel is to be accessible to DOCSIS 1.x Cable Modems, the CMTS MUST follow the DOCSIS 1.x requirements for timing units and relationships for that UCD.

7.1.4.1.2 Mini-Slot Numbering

The MAP counts mini-slots in a 32-bit counter that normally counts to $(2^{(26-M)} - 1)$ and then wraps back to zero. The CMTS MUST match the least-significant bits (i.e. bit 0 to bit 25-M) of the mini-slot counter to 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 mini-slot).

NOTE: The unused upper bits of the 32-bit mini-slot counter (i.e. bit 26-M to bit 31) are unused and MUST be ignored by the CM.

7.1.4.2 S-CDMA Timing Units and Relationships

7.1.4.2.1 Mini-Slot Capacity

On S-CDMA channels, the size of the mini-slot is dependent on the modulation rate, the codes per mini-slot and the spreading intervals per frame, which are all carried in the Upstream Channel Descriptor. The timing units and relationships for S-CDMA are covered in detail in [12]. An example of the timing relationships (assuming 64-QAM modulation) is shown in table 7-2.

Table 7-2: Example of Mini-Slot Capacity in S-CDMA mode

Parameter	Example Value
Spreading intervals per frame	10
Active code length	128
Codes per mini-slot	4
Mini-slots per frame	32
Symbols per mini-slot	40
Bytes per mini-slot	30 (nominal, when using 64-QAM modulation)
Bits/symbol	6 (assuming 64-QAM)
Symbols/second	5 120 000
Mini-slots/second	128 000
Microseconds/mini-slot	250

Note that for S-CDMA the value of Microseconds/mini-slot in table 7-2, is not equal to the inverse of Mini-slots/second since S-CDMA mini-slots are the same length as the frames and are sent out in parallel.

7.1.4.2.2 Mini-Slot Numbering

Mini-slot numbering in S-CDMA mode is described in detail in [12].

7.2 Upstream Data Transmission

7.2.1 Upstream Bandwidth Allocation

The CMTS allocates bandwidth for one or more upstream channels. Bandwidth allocated to one CM may be allocated across multiple channels upon which the CM can transmit.

An upstream channel is modeled as a stream of mini-slots. The CMTS MUST generate the time reference for identifying these slots. The CMTS MUST also control access to these slots by the cable modems. For example, the CMTS 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 the CM's transmission in the time reference specified. This clause describes the elements of the protocol used in requesting, granting and using upstream bandwidth. The basic mechanism for assigning bandwidth management is the allocation MAP (refer to figure 7-1).

The allocation MAP is a MAC Management Message which is transmitted by the CMTS on the downstream channel and which describes, for some interval, the uses of the upstream mini-slots. A given MAP may describe some slots as grants in which particular CMs may transmit data, other slots as available for contention transmission and other slots as an opportunity for new CMs 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.

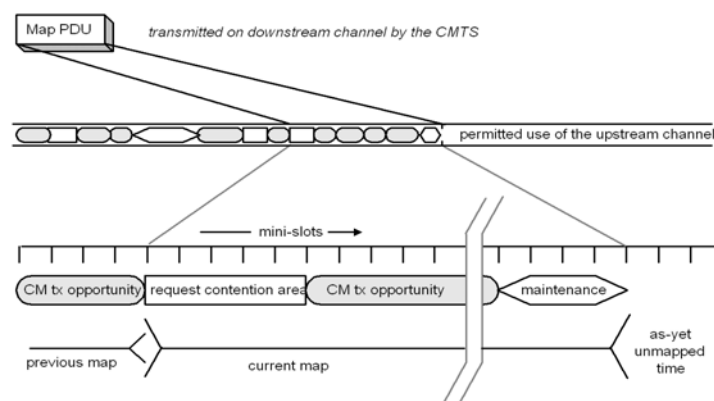


Figure 7-1: Allocation Map

The bandwidth allocation includes the following basic elements:

- Each CM has one or more (14-bit) service identifiers (SIDs) as well as a 48-bit (MAC) address.
- Upstream bandwidth is divided into a stream of mini-slots. Each mini-slot is numbered relative to a master clock reference maintained by the CMTS. The master reference is distributed to the CMs by means of SYNC and UCD messages (see [12]).
- 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 all mini-slots. Mini-slot regions that are not allocated to any transmit opportunities are described by an IE in the MAP assigned to the NULL SID (0x0000). The MAP is described in clause 7.2.1.1.

The CMTS scheduler allocates bandwidth on the individual channels based on the available bandwidth on all of the bonded upstream channels. The CMTS MUST be capable of receiving a request on any channel within the upstream bonding group. The CMTS MUST be capable of granting bandwidth in response to that request on any channel within the upstream bonding group. In this manner, the CMTS MAY dynamically distribute upstream traffic across multiple channels. Similarly, the CMTS MAY consider the physical layer parameters on each of the upstream channels and the requested number of bytes to determine the optimal allocations across channels.

The CMTS uses MAPs to send grants to the CM. Because the upstream parameters of each channel may be very different from each other, the allocation start times of the MAPs may be different from each other as well.

Because the allocation start times and acknowledgment times may vary widely, a CM MUST wait until the acknowledgment time for all upstream channels associated with a given service flow is past the time of request before determining if a re-request is necessary.

CMTSs may ignore part or all of a request. Note that ignoring a request from a CM in Multiple Transmit Channel Mode could result in additional performance degradation (relative to Pre-3.0 DOCSIS) because the CM in Multiple Transmit Channel Mode may take longer to detect lost requests if there are multiple outstanding requests.

7.2.1.1 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 clause 6.4.4. Each IE defines the allowed usage for a range of mini-slots.

NOTE: For TDMA channels, 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 clause 7.1.4.1.2. The relationship between alloc start/ack time counters and the timestamp counter is further described in clause 7.1.4. For DOCSIS 2.0 S-CDMA channels the alloc start/ack time counters are defined in mini-slots which are related to the timestamp counter, frame counter and S-CDMA timestamp snapshot as described in clause 6.4.3.

7.2.1.2 Information Elements

Each IE consists of a 14-bit Service ID (SID), a 4-bit type code (IUC) and a 14-bit starting offset as defined in clause 6.4.4. Since all CMs 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, the CMTS MUST terminate the list with a Null IE (refer to table 6-27).

Five types of Service IDs are defined:

- 1) 0x3FFF - Broadcast, intended for all stations;
- 2) 0x3E00 - 0x3FFE - Multicast, purpose is defined administratively. Refer to annex A;
- 3) 0x2000 - 0x3DFF - Expanded Unicast, intended for a particular CM or a particular service within that CM, when supported by both the CM and CMTS;
- 4) 0x0001 - 0x1FFF - Unicast, intended for a particular CM or a particular service within that CM;
- 5) 0x0000 - Null Address, addressed to no station.

A CM MUST support the Expanded Unicast SID space. A CMTS MAY support the Expanded Unicast SID space.

Unicast SIDs (including Expanded Unicast SIDs) assigned by the CMTS MUST be unique on a given logical upstream. The CMTS MAY support unicast SID assignments which are not unique within a single MAC-sublayer domain as long as they are unique on a given logical upstream.

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.

7.2.1.2.1 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. Clause 7.2.2 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 clause 7.2.3). Packets transmitted in this interval by the CM MUST use either the Request MAC Frame format (refer to clause 6.2.4.3) or the Queue-depth Based Request Format (refer to clause 6.2.4.5).

The Priority Request SIDs are defined in clause A.2.3. These allow contention for Request IEs to be limited to service flows of a given Traffic Priority. (Refer to clause C.2.2.5.1.)

The CMTS **MUST** allocate request opportunities in multiples of the number of mini-slots required to transmit a request on the given channel. For example, if channel one requires 2 mini-slots per request, then the CMTS allocates request regions in multiples of 2 mini-slots. A request region of 5 mini-slots would be illegal on this channel.

7.2.1.2.2 The Request/Data IE

For operation when Multiple Transmit Channel Mode is not enabled, the Request/Data IE provides an upstream interval in which requests for bandwidth or short data packets may be transmitted. For operation when Multiple Transmit Channel Mode is enabled, the Request/Data IE provides an upstream interval in which only requests for bandwidth 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 by the CMTS to specify maximum data length, as well as allowed 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.

This region also supports the Maximum Scheduled Codes feature on Type 3s and Type 4s upstreams as described in [12].

The well-known multicast Service IDs are defined in annex 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 transmitted by a CM in this interval **MUST** indicate in the MAC Header that a data acknowledgment is desired (see table 6-13).

7.2.1.2.3 The Initial Maintenance IE

The Initial Maintenance IE, when used with the Broadcast SID, 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 a Ranging Request (RNG-REQ) message (see clause 7.1.3), **MUST** be provided by the CMTS to allow new stations to perform initial ranging. Packets transmitted by the CM in this interval **MUST** use the RNG-REQ, INIT-RNG-REQ or B-INIT-RNG-REQ MAC Management Message format (refer to clause 7.1.3).

On Type 3 and Type 4 Upstream Channels, the Initial Maintenance IE **MAY** be used by the CM and CMTS with a unicast SID. This is done to provide Unicast Initial Maintenance opportunities in place of Station Maintenance opportunities at the discretion of the CMTS. This may be useful if the first unicast ranging opportunity on an S-CDMA channel needs to have Spreader Off just like initial maintenance, but it is not desirable to impose the overhead of having the Spreader Off on routine Station Maintenance. Unicast Initial Maintenance Opportunities only need to be large enough to allow transmission of the ranging request. The CMTS **MUST NOT** provide unicast Initial Maintenance opportunities on any logical upstream which is not a Type 3 or Type 4 upstream.

7.2.1.2.4 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. The CMTS sends a unicast station maintenance IE to CMs periodically in order to ensure CM upstream transmit signal fidelity. Parameters such as power level, transmit timing, transmit frequency and pre-equalization coefficients can be adjusted in period ranging. Packets transmitted by the CM in this interval **MUST** use the RNG-REQ MAC Management Message format (see clause 6.4.5).

7.2.1.2.5 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 by the CMTS 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).

When Multiple Transmit Channel Mode is not being used, Short Data Grants are used with intervals less than or equal to the maximum burst size for this IUC 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.

With Multiple Transmit Channel Mode, the CM makes requests in number of bytes excluding any physical layer overhead. Therefore, when granting a request, the CMTS assigns a burst profile to the grant. This is indicated by the IUC associated with the IE in the MAP message for the particular grant. When requesting bandwidth while operating in Multiple Transmit Channel Mode, the CM is not constrained by the Maximum Burst Size for Short Data. The CMTS is also not constrained by the Maximum Burst Size for Short Data when granting bandwidth to a CM operating in Multiple Transmit Channel Mode.

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

For Multiple Transmit Channel Mode, the CM MUST be capable of using burst profiles corresponding to Short and Long Data Grants (i.e. IUC 5 and 6) with advanced PHY burst profiles.

7.2.1.2.6 Data Acknowledge IE

The Data Acknowledge IE acknowledges that a data PDU was received. The CMTS MUST include this interval if a CM 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 in a MAP transmitted by the CMTS. This allows cable modems to process all actual interval allocations first, before scanning the MAP for data grants pending and data acknowledgments.

7.2.1.2.7 Expansion IE

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

7.2.1.2.8 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 Data Grant Pending IEs (Data Grants with an inferred length of 0) follow the Null IE.

7.2.1.2.9 Advanced PHY Short and Long Data Grant IEs

These IEs are the Advanced PHY channel equivalent of the Short and Long Data Grant IEs in clause 7.2.1.2.5. In addition, these IEs allow DOCSIS 2.0 modems operating in DOCSIS 2.0 TDMA mode to share the same upstream channel with DOCSIS 1.x modems. Modems registered in DOCSIS 1.x mode MUST NOT use these intervals.

For upstream channels supporting a mixture of DOCSIS 1.x and DOCSIS 2.0 TDMA CMs, the CMTS MUST use the SID in the request and the operational state of the CM to distinguish between requests for IUC 5 and 6 data grants and requests for IUC 9 and 10 data grants (refer to clause 10.2.6). Once this distinction has been made, the CMTS then uses the request size to distinguish between a long grant and a short grant.

Once a CMTS has received a REG-ACK from a 2.0 CM on a Type 2 channel, the CMTS MUST NOT send data grants using IUCs 5 or 6 if either IUC 9 or 10 is defined for that upstream channel. This restriction allows the 2.0 CM to only support 7 burst profiles simultaneously.

With Multiple Transmit Channel Mode, the CM makes requests in number of bytes excluding any physical layer overhead. Therefore, when granting a request, the CMTS assigns a burst profile to the grant. This is indicated by the IUC associated with the IE in the MAP message for the particular grant.

7.2.1.2.10 Advanced PHY Unsolicited Grant IE

This IE can be used by the CMTS to make unsolicited grants of bandwidth to DOCSIS 2.0 CMs. If a significant portion of the traffic for an upstream is going to consist of unsolicited grants of a particular size, this IE provides a way for the CMTS to provide a set of physical layer parameters (such as code word length and FEC length) well tailored to that traffic, without compromising the general usefulness of the Advanced PHY Short or Advanced PHY Long Data Grant IEs. It is never used by the CM to calculate the size of a bandwidth request. The CMTS **MUST NOT** use it to make grants to DOCSIS 1.x CMs.

For Multiple Transmit Channel Mode, the CMTS **MAY** allocate this IE for any data grant. For Multiple Transmit Channel Mode, the CM **MUST** use the burst profile associated with this IE regardless of whether the grant is unsolicited or not.

7.2.1.3 Requesting with Multiple Transmit Channel Mode Disabled

This clause applies to bandwidth requests when Multiple Transmit Channel Mode is disabled, such as when a 3.0 CM is operating on a Pre-3.0 DOCSIS CMTS or a Pre-3.0 DOCSIS CM that does not support Multiple Transmit Channel Mode is operating on a DOCSIS 3.0 CMTS.

Requests refer to the mechanism that a CM uses to indicate to the CMTS that it needs upstream bandwidth allocation. A Request transmitted by a CM **MAY** come as a stand-alone Request Frame transmission (refer to clause 6.2.4.3) or as a piggyback request in the EHDR of another Frame transmission (refer to clause 6.2.5.2).

Request Frames transmitted by a CM **MUST** be sent during one of the following intervals:

- Request IE.
- Request/Data IE.
- Short Data Grant IE*.
- Long Data Grant IE*.
- Adv PHY Short Data Grant IE*.
- Adv PHY Long Data Grant IE*.
- Adv PHY Unsolicited Grant IE*.

NOTE: A request frame could be transmitted during these IEs for the case where Multiple Transmit Channel Mode is disabled for the CM, fragmentation is disabled for a service flow and the CM receives a grant too small to contain the CM's transmission. In this case, the CM may send a request frame in the granted allocation to re-request for the bandwidth.

A piggyback request transmitted by a CM **MUST** be sent in one of the following Extended Headers:

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

A request transmitted by a CM **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 clause 6.2.4.6). The request from the CM **MUST** be large enough to accommodate the entire necessary physical layer overhead (see [12], upstream) 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 clause 6.4.3) or any limits established by QoS parameters associated with the Service Flow.

The CM MUST NOT request more mini-slots than are necessary to transmit the MAC frame. This means that if the CM is using Short and Long Data IUCs to transmit data and the frame can fit into a Short Data Grant, the CM MUST use the Short Data Grant IUC attributes to calculate the amount of bandwidth to request and make a request less than or equal to the Short Data maximum Burst size. If the CM is using Advanced PHY Short and Long Data IUCs to transmit data and the frame can fit into an Advanced PHY Short Data Grant, the CM MUST use the Advanced PHY Short Data Grant IUC attributes to calculate the amount of bandwidth to request and make a request less than or equal to the Advanced PHY Short Data maximum Burst size.

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 that has an ACK Time indicating the request has already been processed until the request is granted or discarded.

7.2.1.4 Requesting with Multiple Transmit Channel Mode Enabled

This clause applies to bandwidth requests when Multiple Transmit Channel Mode is enabled.

As required in clause 6.2.4.3, when the CM is operating in Multiple Transmit Channel Mode, it does not use the Request Frame (bandwidth request in minislots), but rather it uses the Queue-Depth Based Request Frame (bandwidth request in bytes).

Request transmission is controlled by the Request/Transmission Policy parameter on a service flow by service flow basis. For a CM operating in Multiple Transmit Channel Mode, clause 8.3.2 describes that one of the bits of the R/T Policy is used to configure each service flow into one of two modes: Segment Header ON and Segment Header OFF. The requirements for request transmission for a service flow depend on which of these two modes is selected.

7.2.1.4.1 Request Mechanisms for Segment Header OFF Service Flows

As described in clause 8.3.2.2, Segment Header Off operation is only defined for service flows that have a scheduling type of UGS or UGS-AD and as indicated in clause 7.2.3, UGS and UGS-AD service flows are required to have an R/T Policy which prohibits the use of contention request opportunities, request/data opportunities and piggyback requests. As such, the only defined request mechanism for Segment Header Off service flows is the Queue-depth Based Request Frame transmission used to restart grants during a period of rtPS for a UGS-AD service flow (clause 7.2.3.3). The CM MUST be capable of sending a Queue-depth Based Request Frame for a UGS-AD service flow with Segment Headers OFF during a unicast Request IE interval. When sending a Queue-depth Based Request Frame for a UGS-AD service flow, the CM MUST set the number of bytes requested to a non-zero value. Since the CMTS is required to provide fixed-size grants based on the UGS Grant Size parameter, the actual number of bytes requested is irrelevant.

Piggyback requesting for CMs in Multiple Transmit Channel mode is only defined for Segment Header ON operation.

7.2.1.4.2 Request Mechanisms for Segment Header ON Service Flows

For a service flow configured for Segment Header ON operation, the CM can send a Request as a stand-alone Queue-depth Based Request Frame transmission (refer to clause 6.2.4.5) or (unless disabled by the R/T Policy) as a piggyback request in a segment header of another Frame transmission (refer to clause 6.2.5).

The CM MUST be capable of sending a Queue-depth Based Request Frame to request bandwidth for a Segment Header ON service flow during both of the following intervals:

- Request IE.
- Request/Data IE.

A Queue-depth Based Request Frame transmitted by a CM MUST include:

- The Service ID making the request.
- The number of bytes requested with respect to the request byte multiplier for that service flow.

Piggyback requests for a Segment Header ON service flow transmitted by a CM MUST only be sent in the Segment Header Request field of the Segment Header.

A piggyback request transmitted in the Segment Header Request field by a CM MUST include:

- SID Cluster ID associated with the request.
- The number of bytes requested with respect to the request byte multiplier for that service flow.

The CM MUST NOT make a request that would violate limits established by QoS parameters associated with the Service Flow.

The CM MUST NOT request more bytes than are necessary (other than additional bytes required due to the request multiplier) to transmit the data currently queued. In the case where a previous outstanding request was rounded up due to the request multiplier, the CM is not required to decrement a new request by the previous round up amount.

The CM MAY have multiple requests outstanding at a time per Service Flow. If the CMTS does not immediately respond with a Data Grant, the CM is able to unambiguously determine whether its request is still pending by examining MAP messages as discussed in clause 7.2.1.4.2.1.

7.2.1.4.2.1 Queue-Depth Based Request Mechanisms

One mechanism for requesting more upstream bandwidth is to allow the cable modem to request for all the upstream bandwidth it currently needs based on the packets it has ready for upstream transmission. This scheme allows the modem to send up a request based on queue depth where the queue would include all upstream packets and their known MAC headers. This mechanism requires the Continuous Concatenation and Fragmentation feature (discussed in clause 7.2.4) because the CMTS does not know the individual packet boundaries and cannot grant fractions of the request without inadvertently crossing packet boundaries.

When requesting for queue depth, the CM takes into account all packets it wants to transmit and the amount of bandwidth required. This amount of bandwidth includes all known MAC-layer overhead. With Continuous Concatenation and Fragmentation, the CM does not know how many segments the CMTS may use to fragment the grant. For this reason, the CM's bandwidth requests MUST NOT include any estimation for segment headers. The CMTS MUST add the necessary additional bandwidth to compensate for the segment headers when it sends the grant. This is similar to the bandwidth adjustment the CMTS makes when using multiple grant mode of Pre-3.0 DOCSIS Fragmentation.

The CM sends the request for the bandwidth needed for a given service flow on any upstream channel available to the service flow. The CMTS can choose to grant the bandwidth on the upstream channel upon which it received the request, on any other upstream channel associated with the service flow or on any combination of channels associated with the service flow.

In order to provide maximum flexibility in SID assignment on upstream channels, a new term, SID Cluster, is used to define a group of SIDs that contains one SID for each upstream channel associated with a particular Service Flow that is treated the same from a request/grant perspective. An example SID Cluster is shown in table 7-3.

Table 7-3: Example SID Cluster

SID Cluster	US#1 SID	US#2 SID	US#3 SID	US#4 SID
Cluster_0	58	479	85	1 001

A SID Cluster is assigned to a specific service flow on a CM. Whenever the service flow uses a SID Cluster to make a request and a SID is included in the request, the CM MUST use the SID appropriate for the upstream channel on which it is transmitting the request. In the example configuration above, the CM would use SID 479 when sending a bandwidth request on upstream #2. Similarly, whenever the CMTS grants a request that is part of a SID Cluster, it MUST grant the request using the SID corresponding to that SID Cluster on the selected upstream channel. In the example given earlier, if the CMTS chose to use US#3 to grant the request from SID 479 on US#2, the CMTS would place a grant to SID 85 in the MAP for US#3.

The CMTS sends grants spread across channels using individual MAPs for each channel. Should the CMTS decide not to grant all of the bandwidth requested, the CMTS sends a grant-pending in the MAPs for at least one channel until all received requests for that SID Cluster are fulfilled. This is similar to multi-grant mode fragmentation in DOCSIS 1.1. More specifically, when a CMTS issues a grant pending to a CM, the CMTS MUST continue to issue a Data Grant Pending in each MAP on at least one upstream channel associated with the requesting service flow of the CM that has an ACK Time later than the time of the request until the request is granted or discarded. If a CMTS is issuing a Data Grant Pending for a request by a CM on a channel different than the channel on which the request was made, the CMTS MUST include the Data Grant Pending beginning with the first MAP on the channel with an ACK time greater than the translated time of the request.

NOTE: The CMTS may send Data Grant Pending in MAPs prior to those with ACK times greater than the translated time of the request.

In translating the time of the request made on one channel to the time on other channels, the CMTS MUST perform the translation such that the mini-slot count on another channel begins at the exact same time as the beginning mini-slot of the request on the channel on which it was made and if there is no mini-slot that begins at the exact time, the next earliest mini-slot on the other channel is selected. For example, when the CM makes a request on a channel at time T corresponding to minislot M_0 , for each other channel i , the CMTS translates M_0 to the mini-slot count that begins exactly at the time T and if there is no mini-slot beginning exactly at time T , the CMTS translates to M_i equal to the mini-slot count of the next earliest mini-slot on channel i that begins before time T .

Alternatively, the CMTS may choose not to send grant-pendings and allow the CM to re-request for the remainder of the needed bandwidth. This method is similar to the piggyback mode of fragmentation in DOCSIS 1.1. Note that the piggyback mode can add significant latency compared to operation using the multi-grant mode.

The CMTS MUST base ACK times in a MAP on requests originally received on the channel associated with the MAP and not other channels, even if the grants are made on different channels than the channel on which the requests were received. In the absence of received upstream requests or transmissions, the ACK time still needs to be moved forward in time to ensure that CMs can learn the status of outstanding requests that may have been lost.

When the CM makes a request, it MUST remember the mini-slot count on the requesting channel and the mini-slot count on all other channels within the bonding group that starts at the exact time of the request on the requesting channel and if there is no mini-slot that begins at the exact time, then the next later mini-slot count is remembered. For example, when the CM makes a request on a channel at time T corresponding to minislot M_0 , for each other channel i the CM remembers the mini-slot count that begins exactly at the time T and if there is no mini-slot beginning exactly at time T , the CM remembers M_i equal to the mini-slot count of the next later mini-slot on channel i that begins after time T . The CM MUST look for grants to the requesting SID Cluster on all channels associated with the Service Flow. If the acknowledgment time in the MAPs for all channels associated with the Service Flow exceed the time of the request and no grant pendings for the requesting SID Cluster are present in any of those same MAPs, the CM MUST re-request for any ungranted portion of the original request(s). When the CM makes this re-request, it MAY include in the request bandwidth for any new packets requiring transmission.

A CM is allowed to have multiple outstanding requests for a given SID Cluster and may have more than one SID Cluster assigned to the service flow when the service flow is provisioned. Once the CM transmits a request for a service flow, the request/transmission policy for that flow controls whether the CM can make another request for that flow prior to receiving an acknowledgement in the form of a grant or grant-pending. If the request/transmission policy prohibits multiple outstanding requests in contention, the CM MUST NOT request additional bandwidth in contention until all outstanding requests have been granted or expired. The CM MAY piggyback requests for additional bandwidth, even though the CMTS has not fulfilled all previous requests. For example, the CM requests 16 Kbytes in its initial request. The CMTS decides to grant the CM's request with 2 sets of grants of 8 Kbytes each plus segment overhead with the two sets of grants being spaced out in time and appearing in separate MAPs. Once the CM receives the first grant, it may now piggyback request for any new packets that have arrived since the CM made the original request. If the request/transmission policy allows multiple outstanding requests in contention, the CM MAY use contention opportunities to request bandwidth for new packets at any time.

When multiple outstanding contention requests are allowed for a service flow and an additional outstanding contention request is made (see clause 7.2.2.1, regarding collision resolution and contention backoff) the CM MUST increment backoff exponent counts on each channel associated with the service flow (unless the value of Data Backoff End in the MAP message for an upstream channel has already been reached).

When multiple outstanding contention requests are allowed for a service flow, the CM MUST consider a re-request in a contention opportunity due to a previously lost contention request as a retry of a previous request. In other words, the count of request attempts is incremented in this case and results in an increase in the size of the backoff window unless the value of Data Backoff End in the MAP message has already been reached on all upstream channels associated with the service flow (see clause 7.2.2.1, regarding collision resolution and contention backoff). This applies even if additional bandwidth is being requested in the re-request.

More than one SID Cluster may be assigned to a service flow. The CMTS MUST always grant or send grant pendings using the same SID Cluster as the request. The CM MUST stop requesting on a given SID Cluster and switch to another SID Cluster when any one of the following limits is reached (see annex C for details on the TLVs):

- 1) Maximum Requests per SID Cluster - This is the maximum number of requests that can be made using the SID Cluster. Both new requests and re-requests, even for the same bandwidth, increment the count of the number of requests made.
- 2) Maximum Outstanding Bytes per SID Cluster - This is the total size, in bytes, for which there can be outstanding requests using the SID Cluster. Requests for previously unrequested bandwidth increase the outstanding byte count by the total request size, while re-requests increase the count by only the number of newly requested bytes. Grants received for the SID Cluster decrease the count. This is a soft limit, which means that the last request may push the count over the limit, but once the limit has been exceeded, no more requests may be made on this SID Cluster until the SID Cluster has been cleared (all outstanding requested bytes have been granted or outstanding requests have timed out) and operation has switched back to this SID Cluster.
- 3) Maximum Total Bytes Requested per SID Cluster - This is the total number of bytes that can be requested using the SID Cluster. Requests for previously unrequested bandwidth increase the total byte count by the entire request size, while re-requests increase the count by only the number of newly requested bytes. This is a soft limit, which means that the last request may push the count over the limit, but once the limit has been exceeded, no more requests may be made on this SID Cluster until the SID Cluster has been cleared (all outstanding requested bytes have been granted or outstanding requests have timed out) and operation has switched back to this SID Cluster.
- 4) Maximum Time in the SID Cluster - This is the total time, in milliseconds, that a service flow may continue to use the SID Cluster for requests. The start time is initialized to 0 at the time of the first request and is checked before each subsequent request. It should be noted that the final request may actually occur later than this deadline due to the delay between when the limit is checked and when the request is actually made. Once this deadline is reached, no more requests may be made using the SID Cluster.

For all the above SID Cluster switchover criteria, if the service flow has only one SID Cluster and this criterion limit is met, the CM MUST stop making requests and not request again until the SID Cluster has been cleared (any outstanding requested bytes have been granted or outstanding requests have timed out).

The CM MUST NOT request for a given service flow by using more than one SID Cluster at a time. The CM may switch to a different SID Cluster at any time but is required to stop requesting with the current SID Cluster under the conditions given above. Once a CM has stopped using a particular SID Cluster, the CM MUST NOT use the SID Cluster again for requesting until all remaining requests for that SID Cluster have been satisfied. Should the acknowledgment times exceed the requesting time on all channels within the bonding group and there are no grant pendings present in the current MAPs and if the request is still unfulfilled, the CM re-requests for any ungranted bandwidth on that SID Cluster using any of the SID Clusters available for requesting. When switching to a new SID Cluster, the counts corresponding to the first three limits are initialized to 0. When switching to a new SID Cluster, the count corresponding to the Maximum Time in the SID Cluster is set to zero at the time of the first request with the new SID Cluster.

Because the CMTS may use multiple sets of grants to grant the bandwidth from a single request, situations may arise where the CM and CMTS get temporarily out of alignment as requests are lost due to upstream burst errors and collisions and MAPs are lost due to downstream errors. Similar to Pre-3.0 DOCSIS systems, the CM MUST use the acknowledgment time of the requests to decide if the CMTS should have received its request before the CM decides to re-request. Whenever the CM receives a grant-pending for the requesting SID Cluster in the MAP on any channel within the upstream bonding group, the CM MUST NOT re-request for bandwidth for this SID Cluster. Depending on the Request/Transmission Policy Parameters for the service flow, the CM MAY be able to request for new bandwidth ready for upstream transmission for the service flow. Once the CM receives MAPs on all channels within the bonding group with the MAPs containing no grant-pendings for a given SID Cluster and depending on the Request/Transmission Policy Parameters, the CM MAY re-request using piggyback opportunities or contention opportunities for any untransmitted packets whose request time is earlier than the acknowledgment time in the current MAPs. Note that requests whose request time is later than the acknowledgment time may still be in-transit or awaiting processing by the CMTS. The CM MUST wait for the acknowledgment time to be past the requesting time on all channels, within the bonding group, before determining if a re-request is needed. This requirement allows independent operation of CMTS upstream channel schedulers.

As an example of operation during a lost MAP, consider a CM sending a request for 16 Kbytes in its initial request. The CMTS receives the request and sends a set of MAPs (one MAP message for each upstream channel) containing a set of grants for that CM. One of the MAPs is errored due to burst noise so the CM discards the MAP message. Meanwhile, the CM receives unerrored MAPs for the other upstream channels. The CM transmits according to the grants in the correctly-received MAPs. Because the CM has not received a MAP for one of the channels with that MAP containing an acknowledgment time past the time of request, the CM is unable at this point to determine if all of its requests will be granted. The next set of MAPs arrives and the CM sees that the acknowledgment time on all channels is past the time of request and there are no grant pendings for the requesting SID Cluster. The CM knows from this that the CMTS has no outstanding requests for this SID Cluster. However, the CM still has data remaining to be sent from the original 16 Kbyte request. The CM sends a new request for the remainder of the 16 Kbytes plus any new traffic that is ready to be sent upstream for that service flow.

A potential error condition can occur where the CM stops receiving MAPs for one or more upstream channels but continues receiving MAPs for other channels within a service flow's bonding group. The period of time not covered by MAP elements for a channel is considered by the CM as the "unmapped" time for that channel. If the unmapped time on a channel exceeds 1 second, the CM MUST ignore the request time for outstanding requests on that channel (for the purposes of re-requesting) until the CM once again receives MAPs for that channel. This allows the CM to continue requesting for bandwidth on the other channels within a service flow's bonding group when the CM stops receiving MAPs for just some of the channels within the bonding group.

As an example, consider the case where the CM transmits a request for Service Flow A and the time of that request is mini-slot 100 on upstream channel 1, mini-slot 250 on upstream channel 2 and mini-slot 175 on upstream channel 3. Before the CM's request is fully granted, the CM stops receiving MAPs for channel 3 but continues receiving MAPs on channels 1 and 2. The last MAP received for channel 3 had an acknowledgment time of 100. The CM detects that the unmapped time on channel 3 has exceeded 1 second, that it still has not received any MAPs for channel 3 and that the request has not yet been fully granted. The last MAPs received for channels 1 and 2 had acknowledgment times of 790 and 900 respectively. The CM now re-requests for the ungranted portion of the request.

7.2.1.4.2.2 Piggyback Requesting

Piggyback Requesting refers to the use of an extended header of a unicast data transmission for requesting additional bandwidth. The request in effect "piggybacks" on top of a data transmission.

Piggyback requesting is controlled by a bit in the R/T Policy parameter for each upstream service flow (see clause C.2.2.6.3).

Piggyback requesting is performed on a per-service flow basis such that the CM can only piggyback a request for bandwidth on the same service flow for which it is transmitting data.

When a grant pending for one of the CM's SID Clusters occurs in the MAP for any channel within the upstream bonding group, for the service flow associated with that SID Cluster the CM MUST NOT request bandwidth for packets for which the CM previously sent requests using this SID Cluster. The CM MAY piggyback request for packets for which it has not previously sent a request using this SID Cluster or for packets that were requested on another SID Cluster and can be re-requested on this new SID Cluster (per the criteria for re-requesting in clause 7.2.1.4.2.1).

When the CM receives a MAP without a grant pending for the requesting SID Cluster for every channel within the upstream bonding group, the CM MAY re-request for previously requested bandwidth where the request time is earlier than the acknowledgment time in the MAP for all channels within the bonding group. The CM MAY also include in this request the bandwidth for any newly arrived packets.

7.2.1.5 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.

When operating in Multiple Transmit Channel Mode, the CM MUST be able to use the burst profile indicated by the transmission opportunity's IUC, which can correspond to IUC = 1, 2, 3, 4, 5, 6, 9, 10 or 11. This implies that the CM MUST be capable of simultaneously storing nine burst profiles per upstream.

Table 7-4: IE Feature Compatibility Summary for Multiple Transmit Channel Mode

Information Element	Transmit Request Frame	Transmit RNG-REQ	Transmit Any other MAC Frame
Request IE	MUST	MUST NOT	MUST NOT
Request/Data IE	MUST	MUST NOT	MUST NOT
Initial Maintenance IE	MUST NOT	MUST	MUST NOT
Station Maintenance IE	MUST NOT	MUST	MUST NOT
Short Data Grant IE	MAY (Segment HDR OFF only)	MUST NOT	MUST
Long Data Grant IE	MAY (Segment HDR OFF only)	MUST NOT	MUST
Adv PHY Short Data Grant IE	MAY (Segment HDR OFF only)	MUST NOT	MUST
Adv PHY Long Data Grant IE	MAY (Segment HDR OFF only)	MUST NOT	MUST
Adv PHY Unsolicited Grant IE	MAY (Segment HDR OFF only)	MUST NOT	MUST

Table 7-5: IE Feature Compatibility Summary for Pre-3.0 DOCSIS Operation

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
Adv PHY Short Data Grant IE	MAY	MUST	MUST	MUST NOT	MUST
Adv PHY Long Data Grant IE	MAY	MUST	MUST	MUST NOT	MUST
Adv PHY Unsolicited Grant IE	MAY	MUST	MUST	MUST NOT	MUST

7.2.1.6 Map Transmission and Timing

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

- Worst-case round-trip propagation delay - may be network-specific, but on the order of hundreds of microseconds.
- Queuing delays within the CMTS - implementation-specific.
- Processing delays within the CMs - the CMTS MUST allow a minimum processing time by each CM as specified in annex B (CM MAP Processing Time) which has to include any upstream delays caused by upstream interleaving or S-CDMA framing.

- Downstream delays caused by the PMD-layer framer and the FEC interleaver.

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

The CMTS MAY vary the number of mini-slots described from MAP to MAP. At minimum, a MAP transmitted by a CMTS 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 transmitted by a CMTS MAY stretch to tens of milliseconds. Such a MAP would provide poor upstream latency. CMTS 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 transmitted by a CMTS MUST contain two Information Elements: one to describe an interval and a null IE to terminate the list. At a maximum, a MAP transmitted by a CMTS MUST be bounded by a limit of 240 information elements. MAPs are also bounded in that a MAP transmitted by a CMTS MUST NOT describe more than 4 096 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 for each channel. Even though multiple MAPs may be outstanding, for an upstream channel the sum of the number of mini-slots the MAPs transmitted by a CMTS describe MUST NOT exceed 4 096.

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.

The set of all MAPs transmitted by the CMTS, taken together, MUST describe every mini-slot in the upstream channel, whether there is an allocation of an actual transmission opportunity or whether there is an allocation of idle time. If a CM fails to receive a MAP describing a particular interval, it MUST NOT transmit during that interval.

7.2.1.7 Protocol Example

This clause illustrates the interchange between the CM and the CMTS when the CM has data to transmit (figure 7-2). Although the diagram and description are focused on a single upstream channel, DOCSIS 3.0 operation allows for a request to be made on any of multiple upstream channels and the subsequent grants from the CMTS to be one or more transmission opportunities on one or more upstream channels independent of the channel upon which the request was received. Suppose a given CM has a data PDU available for transmission.

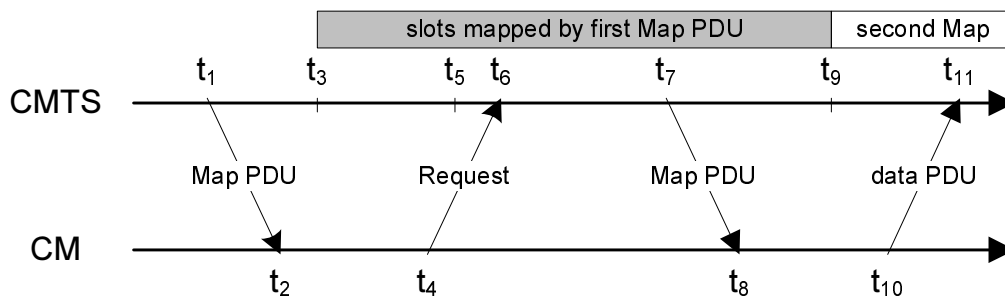


Figure 7-2: Protocol Example

Description steps:

- 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 all the delays discussed in clause 7.2.1.6.
- 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 clause 7.2.2.1, also the multicast SID definitions in clause A.2.2).
- 3) At t_4 , the CM transmits a request for as much bandwidth as needed to accommodate the PDU. Time t_4 is chosen based on the ranging offset (see clause 7.1.3) 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 cannot directly detect the collision. The CM determines that a collision (or other reception failure) occurred when the next MAP with an ACK time indicating that the request would have been received and processed fails to include an acknowledgment of the request. The CM then performs a back-off algorithm and retry (refer to clause 7.2.2.1).

At Step 4, the CMTS scheduler may choose not to accommodate the request within the next MAP. If so, the CMTS MUST reply with a Data Grant Pending in that MAP or discard the request by giving no grant at all. The CMTS MUST continue to send a Data Grant Pending until the request can be granted or is discarded. This signals to the CM that the request is still pending. If the CM is operating in Multiple Transmit Channel Mode and is receiving a Data Grant Pending, it MUST NOT send requests for bandwidth that has already been requested for that service flow. If the CM is not operating in Multiple Transmit Channel Mode and is receiving a Data Grant Pending, it MUST NOT send requests for that service queue.

7.2.1.8 MAP Generation Example - Two Logical Upstreams

This clause illustrates the timing requirements for scheduling an S-CDMA and a TDMA logical channel on the same physical channel.

For simplicity it is assumed that:

- The duration of the S-CDMA frames is an integral multiple of the duration of the TDMA mini-slots.
- Both TDMA and S-CDMA maps begin and end on frame boundaries.
- For the duration of the example there are no S-CDMA bursts with the Spreader Off and there are no Broadcast Initial Ranging regions where both channels are active.

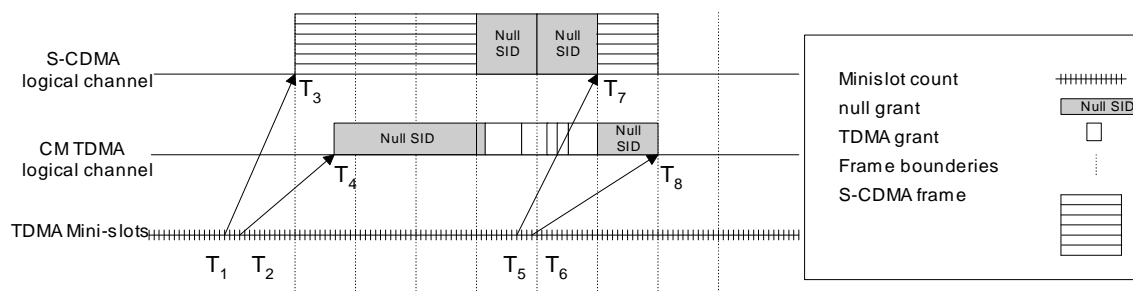


Figure 7-3: Logical S-CDMA TDMA channels

Description:

- 1) The example begins at T_1 and the first MAP discussed takes effect at T_3 .
- 2) At time T_1 , the CMTS transmits a S-CDMA map whose effective starting time is T_3 and end time is T_7 .
- 3) At time T_2 , the CMTS transmits a TDMA map whose effective starting time is T_4 and end time is T_8 .

- 4) At time T_3 the S-CDMA map has three frames of S-CDMA grants. The CMTS upstream scheduler must not allow TDMA transmissions to occur at the same time. To prevent the two channels from interfering with each other the scheduler will mute the TDMA upstream (by granting mini-slots to the NULL SID for the TDMA channel) during the time S-CDMA is active.
- 5) At time T_4 , on a frame boundary, the TDMA channel becomes active. In this example it has one empty mini-slot (NULL SID) to guarantee sufficient guard time for the following TDMA burst. Then it proceeds with usable TDMA grants. At the same time the S-CDMA upstream is muted by granting mini-slots to the NULL SID in every frame.
- 6) At T_5 and T_6 the TDMA logical channel and S-CDMA logical channel transmit the next map for the upstream. Note that figure 7-3 does not continue to detail the complete maps beginning at T_7 and T_8 .
- 7) At time T_7 the S-CDMA map sends a group of S-CDMA grants in a frame.

NOTE: When switching from TDMA to S-CDMA there is no requirement for additional guard time.

7.2.2 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 provide broadcast/multicast request opportunities, which may be subject to collision.

This clause 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 or SID Cluster(s)) it makes these decisions on a per Service Flow or per SID Cluster basis. Refer to annex N for a state transition diagram and more detail.

7.2.2.1 Contention Resolution Overview

The mandatory method of contention resolution which MUST be supported by the CM 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. These values 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 1 023.

For Multiple Transmit Channel Mode, the back-off window values are calculated from the MAPs of the individual channels over which a service flow can be sent.

7.2.2.1.1 Contention Resolution with Multiple Transmit Channel Mode Disabled

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.

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

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 that each IE can represent multiple transmission opportunities.

As an example, consider a CM whose Data Backoff Start is 4 (resulting in an initial back-off window of 0 to 15) and it randomly selects the number 11. The CM defers a total of 11 contention transmission opportunities. If the first available Request IE is for 6 requests, the CM does not use these request opportunities 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 defers for 3 more request opportunities and transmits on the fourth request opportunity within this Request IE.

After a contention transmission, the CM waits for a Data Grant, Data Grant Pending or Data Acknowledge in a subsequent MAP. Once one of these 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. 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.

NOTE 2: Data Acknowledge IEs are intended for collision detection only and 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 not retransmit the data packet. This prevents the CM from sending duplicate packets unnecessarily.

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

NOTE 3: 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 for the calculation of transmit opportunities which is described in clause 7.2.2.2.

7.2.2.1.2 Contention Resolution with Multiple Transmit Channel Mode Enabled

Contention bandwidth requesting in Multiple Transmit Channel Mode is similar to that described above. However, for Multiple Transmit Channel Mode, whenever a service flow is associated with more than one upstream channel, the CM counts the request opportunities in time-order across channels associated with the service flow.

For Multiple Transmit Channel Mode, the CM MUST defer request opportunities across all channels associated with the service flow according to the following rules:

- 1) The CM maintains a data contention backoff window for every service flow. When a switch of SID Cluster occurs, the CM retains current backoff parameter state for each channel over which the service flow operates.
- 2) When the CM initiates a contention request for a bonded service flow, it computes the sum N of the backoff windows defined by the MAPs for all upstream channels associated with the service flow. The backoff window on each channel is equivalent to $2^{\text{Data_Backoff_Start}} - 1$. The CM randomly selects an integer in the range from 0 to N multiplied by the Multiplier to Contention Request Backoff Window (clause C.2.2.6.11) for the service flow.
- 3) The CM orders contention request opportunities across the channels associated with the service flow in time order and transmits its contention request for the service flow after deferring the computed number of opportunities. Whenever the start times of request opportunities on two or more upstream channels align, the CM can pick the ordering of these opportunities as long as all opportunities are counted against the CM's randomly selected backoff value.
- 4) After a contention transmission, the CM waits for a Data Grant or Data Grant Pending in a subsequent MAP. Once either is received, the contention resolution is complete for the case when the CM is not allowed to send requests in contention when there are requests outstanding. The CM determines that the contention transmission was lost when all MAPs for the upstream channels over which the service flow operates do not have a Data Grant or Data Grant Pending for the requesting SID Cluster and have an Ack time more recent than the time of transmission.

- 5) When the CM determines from the MAPs that a contention request was lost, the CM increments the exponent count by one for each of the upstream channels associated with the service flow provided that the Data Backoff End limit has not been reached. If the exponent count already had reached Data Backoff End on a particular channel, then the exponent is not incremented. The CM calculates the sum of the backoff windows over all the channels, performs the backoff as in Rule 2 and transmits the request using the randomly selected opportunity.
- 6) As long as the contention has not been resolved, this retry process continues until the maximum number of consecutive contention retries (16) has been reached, at which time the CM discards from the head of the upstream transmit queue those packets corresponding to the last request transmitted for the service flow. The maximum number of retries is independent of the initial and maximum back-off windows that are defined by the CMTS. When counting request retries for modifying the backoff windows, the CM MUST only count requests sent in contention regions. Thus, in the case that only one outstanding contention request is allowed for the service flow, requests sent in piggyback mode and lost due to noise will not impact the backoff window used by the CM for sending contention requests. In the case of multiple outstanding contention requests, the CM may not know which requests were lost and which were not. So when it is not clear whether a contention request versus piggyback request was lost, the CM MUST assume that a contention request was lost, which will impact the backoff window for the next contention request.

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

While a CM is still attempting to resolve contention for a particular request, the CM MUST ignore changes in values of backoff parameters in MAP messages associated with upstream channels used by the service flow. For any new request that is not a retry, the CM MUST use the backoff parameters in the most recently received MAPs in computing the sum of the backoff windows.

If the CMTS permits multiple outstanding contention requests for the service flow, the CM may transmit additional contention requests. The SID associated with the request can be the same SID or different SID of the service flow's SID Cluster depending on the channel upon which the request is made.

For Multiple Transmit Channel Mode, the CM MUST order request opportunities across all channels associated with the service flow in time order according to the following rules:

- 1) Whenever the start times of TDMA request opportunities on two or more upstream channels are identical, the CM may select the ordering among these opportunities.
- 2) When the channels associated with a service flow have burst profiles that employ upstream interleaving with different latencies or there are some channels that do employ interleaving and others that do not, the CM may select an ordering that reflects when bytes are presented to the physical layer instead of when the request burst is transmitted.
- 3) Whenever multiple contention request opportunities are located in the same S-CDMA frame or when multiple contention request opportunities are located in overlapping S-CDMA frames that are on two or more upstream channels, the CM may select the ordering among these opportunities.
- 4) When different channels have different S-CDMA frame sizes (in symbols), the CM may select an ordering that reflects when bytes are presented to the PHY instead of when the request burst is transmitted.
- 5) If S-CDMA frames containing contention opportunities overlap in time with TDMA contention opportunities on other channels, the CM may select the ordering of these opportunities. In this specific case, an additional allowance is provided for the TDMA contention opportunities in relation to the S-CDMA opportunities on other channels whereby the CM can select how a TDMA opportunity is ordered with respect to S-CDMA contention opportunities in the S-CDMA frame that overlaps the TDMA opportunity or the frame just before or the frame just after.
- 6) TDMA contention opportunities on a channel shall be deferred in time order, although not necessarily consecutively due to opportunities on other channels.
- 7) S-CDMA contention opportunities in a later S-CDMA frame shall not be ordered prior to contention opportunities in an earlier S-CDMA frame on the same channel.

A CM transmitting a RNG-REQ, B-INIT-RNG-REQ or INIT-RNG-REQ in the Initial Maintenance IE MUST perform truncated binary exponential backoff on the single channel itself using the Ranging Backoff Start and Ranging Backoff End of the MAP associated with the channel to control the backoff window. Contention resolution on a single channel is performed as described in the clause applying to Pre-3.0 DOCSIS operation (clause 7.2.2.1.1).

7.2.2.2 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 contention 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, that is, 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 a SID of 0x3FF4 (refer to annex A), then a CM can potentially start a transmission 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 a Broadcast 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, Adv PHY Short/Long Data Grant IEs, Adv PHY Unsolicited Grant IEs, unicast Initial Maintenance 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.

This is summarized in table 7-6.

Table 7-6: Transmit Opportunity Summary

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

NOTE: Transmit Opportunity should not be confused with Burst Size. Burst Size requirements are specified in table 6-26.

For Multiple Transmit Channel Mode, the CM MUST place traffic into segments based on the start time of each segment. Traffic at the head of the service flow queue MUST be placed into the segment which is transmitted first with the following exceptions:

- Whenever the start times of TDMA transmit opportunities on two or more upstream channels are identical, the CM may select the ordering among these opportunities.
- When multiple channels are associated with a service flow and have burst profiles that employ interleaving with different latencies or there are some channels that do employ interleaving and others that do not, the CM may select an ordering that reflects when bytes are presented to the physical layer instead of when the data burst is transmitted.
- Whenever transmit opportunities for a service flow are located in overlapping S-CDMA frames that are on two or more upstream channels, the CM may select the ordering among these opportunities.
- When different channels have different S-CDMA frame sizes (in symbols), the CM may select an ordering that reflects when bytes are presented to the PHY layer instead of when the burst is transmitted.

- If S-CDMA frames containing transmission opportunities for a service flow overlap in time with TDMA transmission opportunities on other channels, the CM may select the ordering of these opportunities. In this specific case, an additional allowance is provided for the TDMA opportunities in relation to the S-CDMA opportunities on other channels whereby the CM can select how a TDMA opportunity is ordered with respect to S-CDMA opportunities in the S-CDMA frame that overlaps the TDMA opportunity or the frame just before or the frame just after.
- TDMA transmit opportunities on a channel shall be used for segmentation in time order.
- S-CDMA transmission opportunities in a later S-CDMA frame shall not be ordered prior to transmission opportunities in an earlier S-CDMA frame on the same channel.

7.2.2.3 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 clause C.2.2.6.3).

- 1) When a CM has data to send, it **MUST** first use any available Data Grants assigned to the Service Flow or Class of Service if it is allowed to do so. If there are no Data Grants, the CM **MUST** then use an available unicast request opportunity. If there are no unicast request opportunities, then the CM can use broadcast/multicast request opportunities for which it is eligible while complying with the contention backoff requirements in clause 7.2.2.1. The intent is that the CM use Data Grants to send data when it is able to do so and if it needs to request, then it looks for a non-contention request, if available, to make a request before resorting to request opportunities in contention. The use of piggybacked requests relative to other types of requests is left unspecified, except for requirements in clause 7.2.5.2.
- 2) For Multiple Transmit Channel Mode, a CM **MUST NOT** have more requests outstanding per SID Cluster than the Maximum Requests per SID Cluster, which is a parameter that may be included in the registration response message. When Multiple Transmit Channel Mode is disabled, a CM **MUST NOT** have more than one Request outstanding at a time for a particular Service ID.
- 3) When Multiple Transmit Channel Mode is disabled, if a CM has a Request outstanding it **MUST NOT** use intervening contention intervals for that Service ID.

7.2.3 Upstream Service Flow Scheduling Services

The following clauses 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 annex C. The clause 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 service. Table 7-7 shows the relationship between the scheduling services and the related QoS parameters.

7.2.3.1 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. UGS 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 setting (refer to clause C.2.2.6.3) 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. The CMTS MUST reject a UGS Service Flow for which the Request/Transmission Policy contains the value zero for any of the bits 0-4. 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 annex O).

The Unsolicited Grant Synchronization Header (UGSH) in the Service Flow EH Element (refer to clause 6.2.5.4.2) 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 by the CM 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:

$$\text{Max}(T) = T * (R * 1.01) + 3B$$

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

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

UGS services can be configured for either segment header-on or segment header-off.

As described in clause 8.3.2.2, the CMTS will generally not allocate bandwidth on more than one upstream channel for a UGS flow with Segment Header OFF. An exception to this might be a UGS flow for which unambiguous grant ordering is enforced by the selection of a Nominal Grant Interval that is less (by some margin) than the Tolerated Grant Jitter. In such a service flow, packet ordering can be assured without the need and overhead of segment headers.

7.2.3.2 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 clause 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 MUST reject an rtPS Service Flow for which the Request/Transmission Policy contains the value zero for any of the bits 0-4. 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.

7.2.3.3 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 UGS/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. The CMTS 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 clause 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. The CMTS MUST reject a UGS-AD Service Flow for which the Request/Transmission Policy contains the value zero for any of the bits 0-4. 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 annex O.) Because the Service Flow is provisioned as a UGS flow with a specific grant interval and grant size, when restarting UGS with MTC mode disabled, the CM MUST NOT request a different sized grant than the already provisioned UGS flow. When MTC mode is enabled, the CM is allowed to send any non-zero value for the request. 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.

7.2.3.4 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 service flows 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 clause C.2.2.6.3) should be such that the CM is allowed to use contention request opportunities. The CMTS MUST reject an nrtPS Service Flow for which the Request/Transmission Policy contains the value one for any of the bits 0-2. 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.

7.2.3.5 Best Effort Service

The intent of the Best Effort 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. The CMTS MUST reject a Best Effort Service Flow for which the Request/Transmission Policy contains the value one for any of the bits 0-2. 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.

7.2.3.6 Other Services

7.2.3.6.1 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.

7.2.3.7 Parameter Applicability for Upstream Service Scheduling

Table 7-7 summarizes the relationship between the scheduling services and key QoS parameters. A detailed description of each QoS parameter is provided in clause C.2.

Table 7-7: Parameter Applicability for Upstream Service Scheduling

Service Flow Parameter	Best Effort	Non-Real-Time Polling	Real-Time Polling	Unsolicited Grant	Unsolicited Grant with Activity Detection
Miscellaneous					
Traffic Priority	Optional Default = 0	Optional Default = 0	N/A (see note 1)	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 = 3 044	Optional Dflt = 3 044	Optional Dflt = 3 044	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 (see note 2)	Optional (see note 3)	Optional (see note 3)	Optional (see note 3)	Optional (see note 3)
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 (see note 3)	Mandatory	N/A	Optional (see note 2)
Tolerated Poll Jitter	N/A	N/A	Optional (see note 3)	N/A	Optional (see note 3)
NOTE 1: 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.					
NOTE 2: Default is same as Nominal Grant Interval.					
NOTE 3: Default is CMTS-specific.					

7.2.3.8 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 clause 7.2.2.3 and the Request/Transmission Policy specified for the service flow.

7.2.4 Continuous Concatenation and Fragmentation

CMs in Multiple Transmit Channel Mode generally use Continuous Concatenation and Fragmentation (CCF) to fill data grants. CCF treats each service flow at the CM as a continuous stream of data regardless of the channel on which that data is transmitted. Each service flow is a different stream. When in Multiple Transmit Channel Mode, the CM MUST NOT use Pre-3.0 DOCSIS concatenation or Pre-3.0 DOCSIS fragmentation for any upstream service flow. When in Multiple Transmit Channel Mode, the CM MUST use CCF for upstream service flows configured for segment-header-on operation. CCF operates on a segment basis where a segment is an individual data grant to a service flow. CCF packs the grants with data in a streaming manner. The segmentation with CCF is performed on a per-service flow basis.

With CCF, a segment header at the beginning of each segment contains a pointer to the beginning of the first MAC header contained in the segment. This is similar to the use of the MPEG pointer for locating the MAC frame boundaries in the downstream. Also contained in the segment header is a sequence number to show where the payload of this segment should be placed at the CMTS relative to payloads for other segments for this service flow. Due to varying propagation delays and overlapping segments on different channels, the segments are not guaranteed to arrive in order at the CMTS MAC. After the segment header, the CM places the next MAC bytes to be transmitted regardless of packet boundaries (there is no concatenation header or fragment header inserted with the data). The CM MUST fill each segment in the order of the rules given in clause 7.2.2. The CM MUST increment the sequence number in the segment header according to the order the CM is filling the segments for the service flow. As long as the CM has upstream traffic for a given service flow, it MUST completely fill each segment with the upstream traffic. Once the CM has partially filled a segment and there is no other MAC data available for transmission for that service flow, the CM MUST pad the remainder of the segment according to the rules specified in the FEC coding portions of the [12].

Figure 7-4 shows an example of CCF with segment headers.

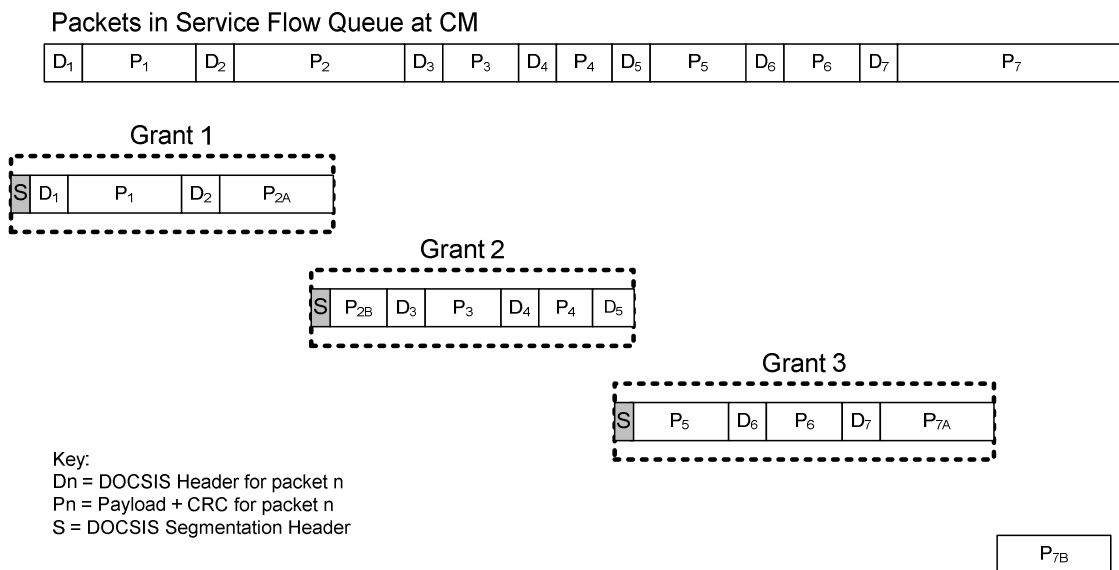


Figure 7-4: CCF using Segment Headers

The pointer field in the segment header allows the CMTS to find the packet boundaries in the event of a lost segment. The pointer in the segment header in grant 1 would point to the first byte after the segment header. The pointer in the segment header in grant 2 would point to the DOCSIS header of packet 3. The pointer in the segment header in grant 3 would point to the DOCSIS header of packet 6. Thus, if any segment is lost, the CMTS can still find the packet boundaries in the remaining segments. The CMTS MAC uses the grant size to determine how many MAC bytes to extract from each grant.

7.2.5 Pre-3.0 DOCSIS Concatenation and Fragmentation

As stated in clause 7.2.4, all upstream service flows on CMs operating in Multiple Transmit Channel Mode will not use Pre-3.0 DOCSIS concatenation or Pre-3.0 DOCSIS fragmentation. Upstream service flows on CMs not operating in Multiple Transmit Channel Mode may use Pre-3.0 DOCSIS Concatenation and Fragmentation as described below.

7.2.5.1 Concatenation

A Specific MAC Header is defined to allow multiple MAC frames to be concatenated when Multiple Transmit Channel Mode is disabled. This allows a single MAC "burst" to be transferred across the network. The PHY overhead and the Concatenation MAC Header only occur once.

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

Concatenation of multiple MAC frames by the CM MUST be as shown in figure 7-5. 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 with Multiple Transmit Channel Mode disabled.

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




Figure 7-5: Concatenation of Multiple MAC Frames

Only one Concatenation MAC Header MUST be present per MAC "burst" transmitted by a CM. The CM MUST NOT use nested concatenation. Immediately following the Concatenation MAC Header transmitted by the CM 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 transmitted by a CM MUST be unique and can be of any type. This means that Packet and MAC-specific Frames MAY be mixed together by a CM in a single concatenation. However, all frames in a concatenation transmitted by a CM MUST be assigned to the same Service Flow. The CMTS MUST support concatenations containing multiple frame types, including both Packet and MAC-specific frames.

The embedded MAC frames can be addressed to different destinations and MUST be delivered by the CMTS as if they were transmitted individually.

7.2.5.2 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. Once fragmentation is enabled for a 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.

Figure 7-6 provides fragmentation details.



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

Fragmentation headers transmitted by the CM 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 by the CM **MUST** match the SID used in the Partial Grant that initiated the fragmentation. A separate CRC **MUST** be calculated by the CM 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.

7.2.5.2.1 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 7-7. The phrase "untransmitted portion of frame" 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.

7.2.5.2.1.1 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, any piggyback request for the next fragment **MUST** be made by the CM in the BP_UP EHDR portion of the fragment header. Any piggyback request for a subsequent frame **SHOULD** be made by the CM in the BP_UP EHDR portion of the last fragment, but can be made in one of the extended headers inside the original frame. However, the same request **MUST NOT** be made by the CM 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.

- 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 by the CM MUST NOT be fragmented.

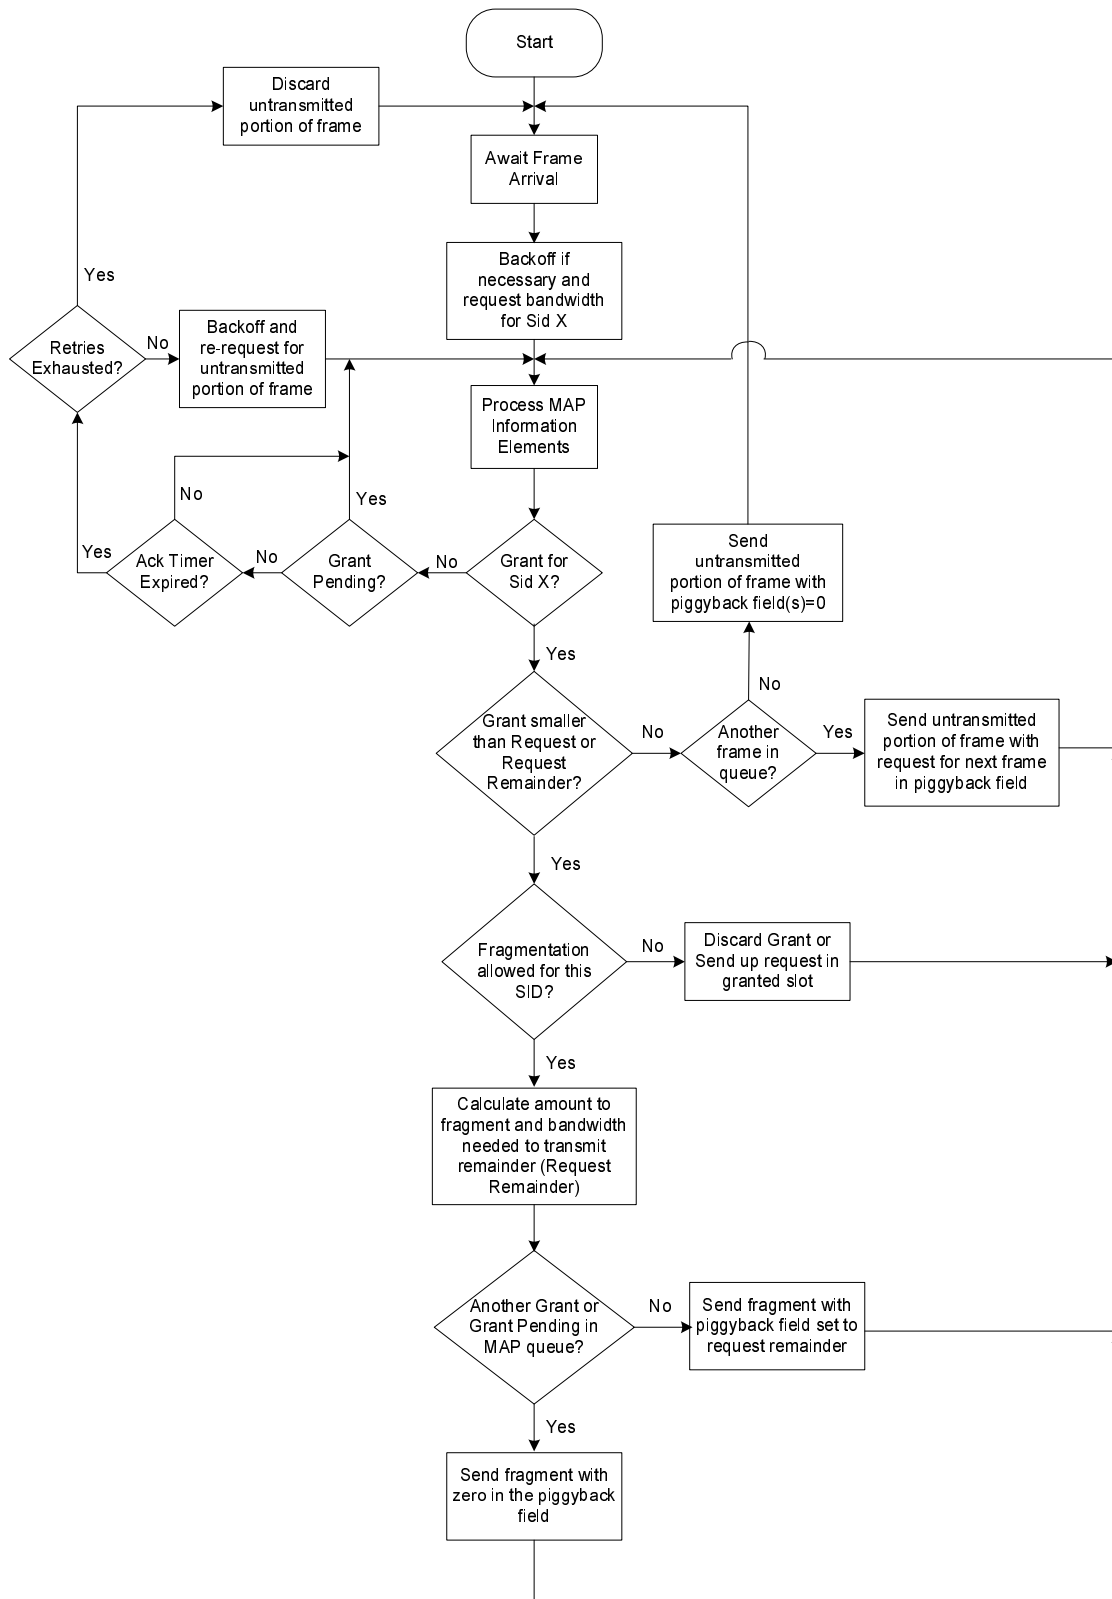


Figure 7-7: CM Fragmentation Flowchart

7.2.5.2.2 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 Piggyback 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.

A CMTS **MUST** support Multiple Grant Mode or Piggyback Mode for performing fragmentation. A CMTS **MAY** support both fragmentation modes.

7.2.5.2.2.1 Multiple Grant Mode

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 miss a MAP with a partial grant and thus request 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.

7.2.5.2.2.2 Piggyback Mode

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 queued 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 by the CMTS, but they may be ignored also.

7.2.5.2.3 Fragmentation Example

7.2.5.2.3.1 Single Packet Fragmentation

Refer to figure 7-8. 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 mini-slots. 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 mini-slots. The CM calculates how much MAC information can be sent in the granted number of mini-slots using the specified burst profile. In the example in figure 7-8, 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 no 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 mini-slots 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 7-8, the CM sends up a request for enough mini-slots to hold the POH plus 150 bytes ($1\ 018 - 884 + 16$). 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 mini-slots 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 mini-slots 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 another grant or a grant pending 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 or grant pending, the CM calculates the number of mini-slots 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 7-8, the grant is large enough to hold the remaining 150 bytes plus POH.
- 11) **Normal operation:** The CM then returns to the normal operation of waiting for grants and requesting for bandwidth. 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.

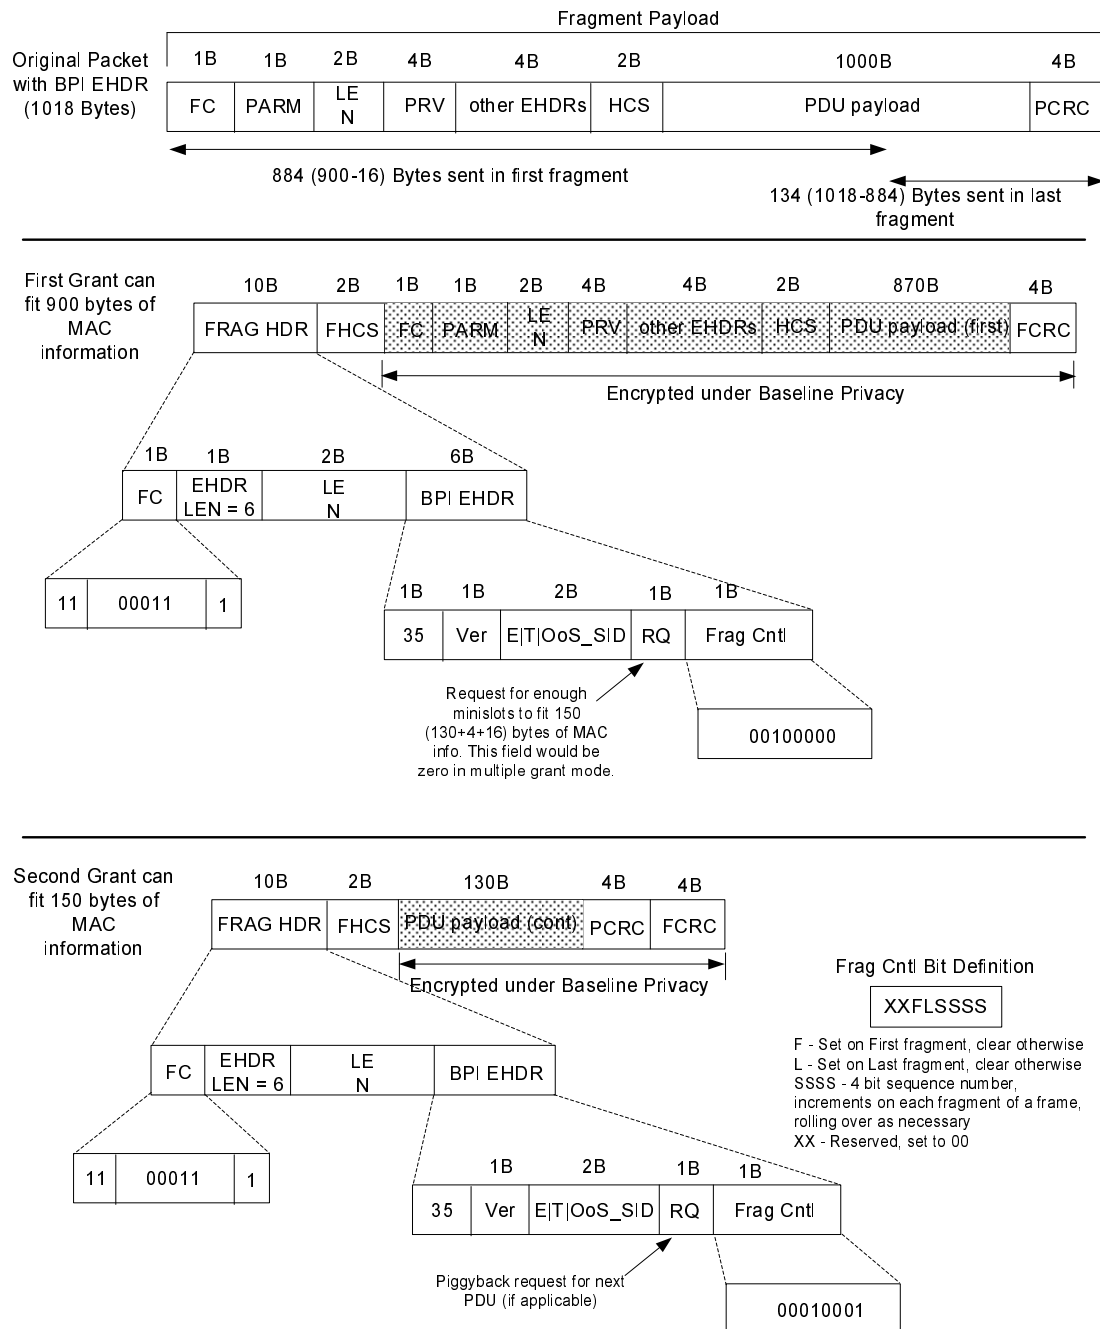


Figure 7-8: Example of Fragmenting a Single Packet

7.2.5.2.3.2 Concatenated Packet Fragmentation

After the CM creates the concatenated packet, the CM treats the concatenated packet as a single PDU. Figure 7-9 shows an example of a concatenated packet broken into 3 fragments.

NOTE: The packet is fragmented without regard to the packet boundaries within the concatenated packet.

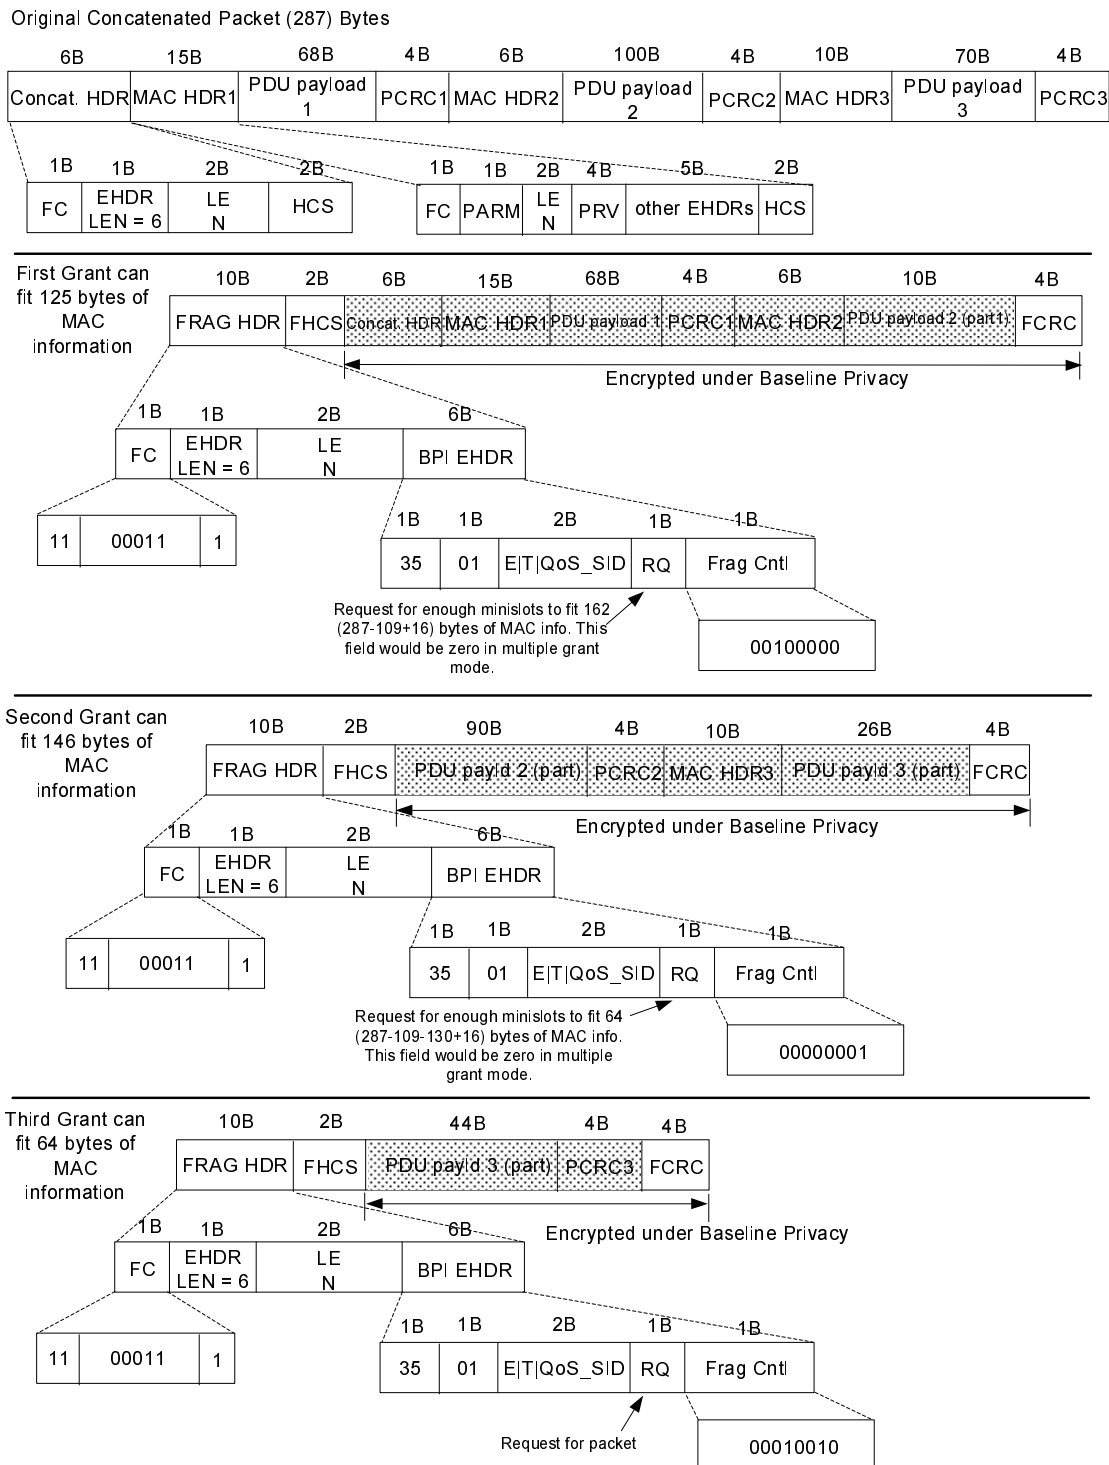


Figure 7-9: Fragmented Concatenated Packet Example

7.2.5.2.4 Pre-Registration Fragmentation

A CM MUST support Pre-3.0 DOCSIS fragmentation prior to registration. A CM signals this capability to the CMTS by setting the appropriate Capability Flags bit in the B-INIT-RNG-REQ as per clause 6.4.5.3.1. A CM may also receive a partial grant prior to registration when MSC is enabled as detailed in [12].

7.2.5.2.5 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 never sent encrypted as a result of being concatenated and fragmented.

7.3 Upstream - Downstream Channel Association within a MAC Domain

7.3.1 Primary Downstream Channels

During initialization, the CM **MUST** select a single downstream channel with which to attempt initial ranging. This downstream channel is known as the CM's candidate Primary Downstream Channel. Over the course of normal operation, the CM's Primary Downstream Channel may be changed several times by the CMTS. The CM receives SYNC messages only on its Primary Downstream Channel. The CM **MUST** ignore SYNC messages received on any downstream channel other than its Primary Downstream Channel. The CM receives and parses the MDD message from its Primary Downstream Channel for information to perform operations including plant topology resolution and initial upstream channel selection (see clause 10.2).

A CM **MUST** change its candidate Primary Downstream Channel in response to a Downstream Frequency Override encoding in a RNG-RSP. A CM **MUST** change its Primary Downstream Channel in response to the Dynamic Bonding Change mechanism as described in clause 11.5.

During initialization, the CM is required to receive only those MAPs and UCDs which are sent on its Primary Downstream Channel. For this reason, it is necessary for the Primary Downstream Channel to carry UCDs and MAPs for the upstream channel(s) upon which the CM will attempt initial ranging. Upon transmission of a REG-ACK message with a confirmation code of success(0), the CM **MUST** be capable of receiving MAPs and UCDs on all of the Downstream Channels in its Receive Channel Set. The CM identifies which Downstream Channels in its Receive Channel Set carry the UCDs and MAPs for the upstream channels in its Transmit Channel Set. In the event that a particular Downstream Channel, upon which the CM is receiving UCDs and MAPs, becomes unavailable (e.g. via DBC or channel failure) or if the CM detects that UCDs and MAPs are no longer available on that channel, the CM **MUST** select a different Downstream Channel in its Receive Channel Set as the source for those UCDs and MAPs, if they are available.

As defined in clause 9.1, the CM forwards broadcast data packets which are non-sequenced and unlabeled (i.e. broadcast data packets that do not have a DSID) that are received on its Primary Downstream Channel and discards any such packets received on a channel other than the CM's Primary Downstream Channel.

For Integrated CMTS implementations with four or fewer downstream channels per RF port, the CMTS **MUST** support configuring all downstream channels to be primary-capable. For Integrated CMTS implementations with greater than four downstream channels per RF port, the CMTS **MUST** support configuring a minimum of 4 downstream channels to be primary-capable. For Modular CMTS implementations, the CMTS **SHOULD** support configuring a sufficient number of downstream channels as primary-capable in order to meet operator deployment needs. The CMTS transmits the following information on each Primary-Capable Downstream Channel:

- SYNC messages.
- MDD messages containing all of the TLVs required for a Primary-Capable Channel per clause 6.4.28.
- UCDs and MAPs for each upstream channel listed in the MDD Upstream Ambiguity Resolution Channel List.

The CMTS does not transmit SYNC messages on downstream channels that are not configured as Primary-Capable.

NOTE: Pre-3.0 DOCSIS CMs are unable to use non-primary-capable downstream channels. As a result, a CMTS is not required to support functionality that is needed for pre-3.0 DOCSIS CMs (such as DES encryption) on these downstream channels.

7.3.2 MAP and UCD Messages

UCD and MAP messages for a given upstream channel may be sent on any downstream channel in the MAC Domain, regardless of whether or not the channel is a Primary-Capable Downstream Channel. UCD and MAP messages for a given upstream channel may be sent on more than one downstream channel. A CMTS MUST transmit MAP and UCD messages for each upstream channel in a CM's Transmit Channel Set on at least one downstream channel in that CM's Receive Channel Set. The CMTS MUST ensure that the UCDs and MAPs for a given upstream channel are identical on all downstream channels on which they are transmitted. Since each CM is only required to receive MAP messages for a particular upstream channel on a single downstream channel, the CMTS MUST transmit all of the MAPs for a given upstream channel on each of the downstream channels on which those MAPs are carried. The CMTS MUST transmit all UCDs for a particular upstream channel on each of the downstream channels on which the MAPs for that upstream channel are transmitted. On each Primary-Capable Downstream Channel, the CMTS MUST transmit UCDs and MAPs for each upstream channel listed in the Upstream Ambiguity Resolution Channel List TLV contained in the MDD on that downstream channel.

7.3.3 Multiple MAC Domains

The CMTS might operate in a configuration in which downstream channels are shared across multiple MAC domains. If a downstream channel is shared between multiple MAC domains, the CMTS MUST ensure that the downstream channel is primary-capable in only one of the MAC domains.

On a given downstream channel, the CMTS MUST ensure that MAPs and UCDs are transmitted for only a single MAC domain. If a downstream channel is primary capable and shared across multiple MAC domains, the CMTS MUST include the MAP and UCD Transport Indicator TLV in the MDD message.

If the MAP and UCD Transport Indicator TLV is present in the MDD message, the CM MUST restrict the set of channels on which it receives MAPs and UCDs to those indicated by the MAP and UCD Transport Indicator TLV. If the MAP and UCD Transport Indicator TLV is not present in the MDD message, the CM can receive MAPs and UCDs from any of the Downstream Channels in its Receive Channel Set per the Primary Downstream Channels clause.

7.4 DSID Definition

A DSID is a 20-bit value contained in a Downstream Service Extended Header (DS EHDR) on a frame that identifies a stream of packets to a set of CMs (ref: DS EHDR clause 6.2.5.6). The CM uses the DSID for purposes of downstream resequencing, filtering and forwarding. A DSID value communicated to the CM by the CMTS is said to be "known" by the CM. Any DSID value not communicated to a CM is considered to be "unknown" by the CM.

The CMTS inserts a Downstream Service Extended Header (DS EHDR) on each sequenced downstream packet to provide the DSID value and the packet's sequence number specific to that DSID. The use of a DSID to identify a particular packet stream sequence allows DOCSIS 3.0 CMs to filter downstream packets based on the DSID value and resequence only those packets intended to be forwarded through the CM.

The CMTS labels all packets of a multicast session with a DSID and communicates that DSID to the set of CMs that are intended to forward that session. DOCSIS 3.0 CMs will only forward multicast traffic that is labeled with a known DSID. In order to reach all the intended recipients, the CMTS replicates a multicast packet as necessary among the downstream channels of a MAC Domain. The CMTS inserts a DS EHDR on multicast packets to provide the DSID which identifies the CM or set of CMs that will forward a particular replication of a multicast session.

A DSID used to provide sequenced delivery of packets and hence to identify a resequencing context in the CM, is termed a Resequencing DSID. A DSID used to label multicast packets is termed a Multicast DSID. A DSID can be used simultaneously for both purposes (e.g. sequenced multicast delivery).

The stream of packets identified by a DSID is independent of a CMTS service flow. For example, the CMTS may transmit packets labeled with the same DSID for one or more Individual Service Flows forwarded to the same CM. Alternatively, the CMTS may classify different IP multicast sessions each with different DSIDs to the same "Group" Service Flow (clause 7.5.8).

A CMTS communicates DSIDs to CMs with the following messages:

- The MDD message contains a "Pre Registration DSID" intended for pre-registration downstream multicasts (see clause 6.4.28.1.5).
- The REG-RSP or REG-RSP-MP message contains DSID Encodings that define an initial set of DSIDs to be recognized by the CM (see clause C.1.5.3.8).
- The DBC-REQ message dynamically updates the set of DSIDs recognized by the CM after registration (add, delete or modify), see clause 11.4.1.2.

The CMTS MUST assign DSID values uniquely per MAC Domain. This simplifies operational reporting of DSIDs by the CMTS. DSID values are intended to be internally assigned by the CMTS and not externally assigned by an OSSI application.

The CM MUST report the total number of DSIDs it supports for filtering purposes (clause C.1.3.1.30). The CM also MUST report the number of Resequencing DSIDs (clause C.1.3.1.31) and the number of Multicast DSIDs supported (clause C.1.3.1.32). The CM MUST support one DSID-Indexed Payload Header Suppression Rule on each Multicast DSID supported. The CM MUST report at least 32 Total DSIDs, 16 Resequencing DSIDs and 16 Multicast DSIDs. If the CM reports values larger than the minimum for any of the DSID capabilities, the Total DSIDs may be less than the sum of the Resequencing DSIDs and Multicast DSIDs to allow for CM optimization of resource utilization.

The CMTS MUST NOT signal a CM to add more DSIDs than the CM reports in the Total Downstream Service ID Support capability encoding (clause C.1.3.1.30). The CMTS MUST NOT signal a CM to add more Resequencing DSIDs than the CM reports in the Resequencing Downstream Service ID Support capability (clause C.1.3.1.31). The CMTS MUST NOT signal a CM to add more Multicast DSIDs than the CM reports in the Multicast Downstream Service ID (DSID) Support capability encoding (clause C.1.3.1.32).

7.5 Quality of Service

7.5.1 Concepts

7.5.1.1 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. 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 mini-slots and the expected behavior of the CMTS upstream scheduler.

A Service Flow is partially characterized by the following attributes.

ServiceFlowID: exists for all service flows.

SID Cluster Group: defines the set of SID Clusters assigned to a service flow. It 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 AuthorizedQoSParamSet allowed by the authorization module. The ProvisionedQoSParamSet is defined once when the Service Flow is created via registration.

AuthorizedQoSParamSet: defines a set of QoS Parameters which define the maximum collection of resources that a particular flow is authorized to use. Any subsequent flow requests will be compared against the AuthorizedQoSParamSet. The AuthorizedQoSParamSet is communicated to the CMTS through a means other than the configuration file.

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.

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 figures 7-10 and 7-11. The ActiveQoSParameterSet is always a subset of the AdmittedQoSParameterSet which is always a subset of the AuthorizedQoSParameterSet. 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 fewer resources (e.g. Maximum Traffic Rate), A is a subset of B if the parameter in A is less than or equal to the same parameter in B.
- If a larger QoS parameter value indicates fewer 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.

In the dynamic authorization model, the authorized envelope (the AuthQoSParamSet) is determined by the Authorization Module. In the provisioned authorization model, the authorized envelope is determined by the ProvisionedQoSParameterSet (refer to clause 7.5.4).

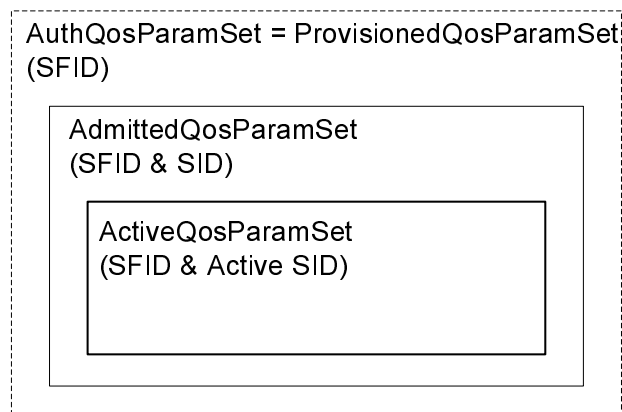


Figure 7-10: Provisioned Authorization Model "Envelopes"

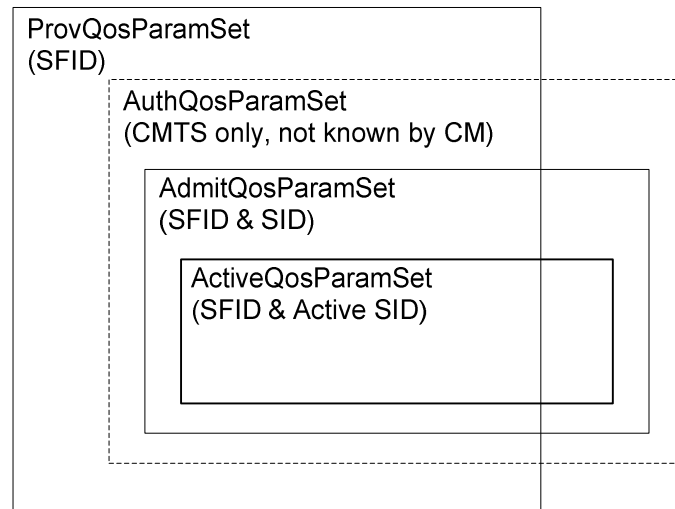


Figure 7-11: Dynamic Authorization Model "Envelopes"

It is useful to think of four states of Service Flows:

Provisioned: A Service Flow in this state is known via provisioning through the configuration file, its AdmittedQoSParamSet and ActiveQoSParamSet are both null.

Authorized: A Service Flow in this state is known to the CMTS via an outside communication mechanism, its AdmittedQoSParamSet and ActiveQoSParamSet are both null. Authorized service flows are not normally communicated to the CM.

Admitted: A Service Flow in this state 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 signaled by some other mechanism. Generally, Admitted Service Flows have associated Classifiers, however, it is possible for Admitted Service Flows to use policy-based classification.

Active: A Service Flow in this state 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.

An inactive service flow may or may not have associated Classifiers. If an inactive service flow has associated Classifiers, the Classifiers MUST NOT be used by a CM or CMTS to classify packets onto the flow, regardless of Classifier Activation State.

7.5.1.2 Classifiers

A Classifier is a set of matching criteria applied to each packet entering the cable network which consists of some packet matching criteria (destination IP address, for example) and a classifier priority. A QoS Classifier additionally consists of a reference to a service flow. If a packet matches the specified packet matching criteria of a QoS Classifier, it is then delivered on the referenced service flow. An Upstream Drop Classifier is a Classifier created by the CM to filter upstream traffic. If a packet matches the specified packet matching criteria of an Upstream Drop Classifier, it is then dropped.

7.5.1.2.1 Upstream and Downstream QoS Classifiers

Several QoS 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 clause 7.5.6.1). 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 7-12 illustrates the mapping discussed above.

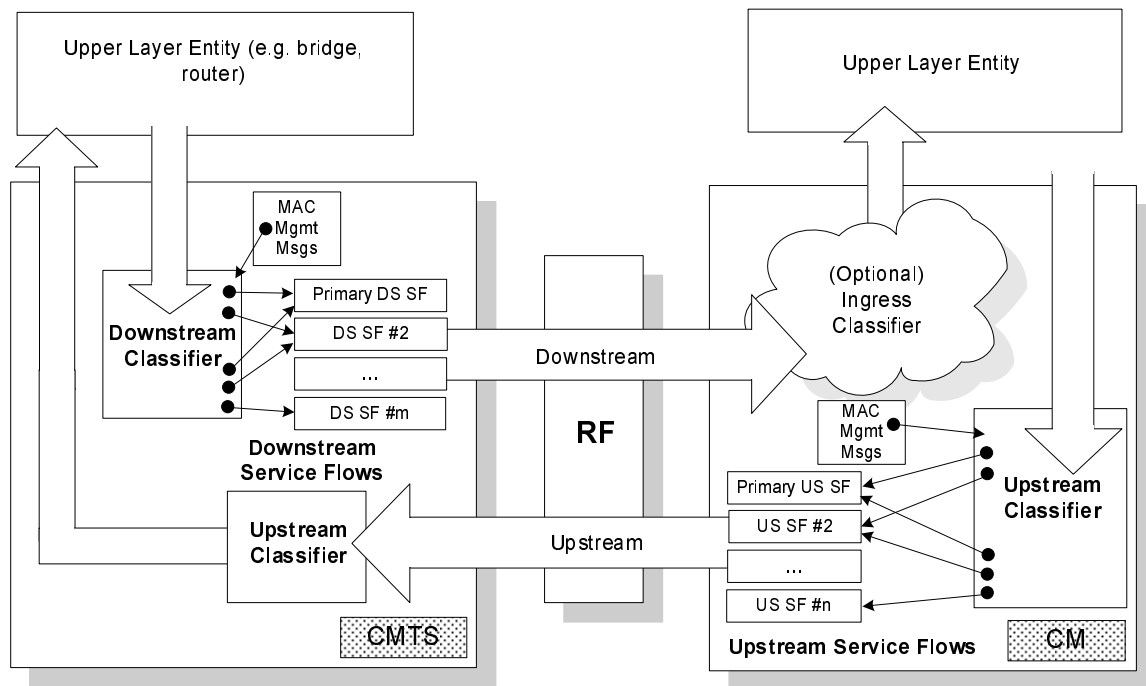


Figure 7-12: Classification within the MAC Layer

The highest priority Classifier MUST be applied by the CM or CMTS first. If a Classifier is found that has at least one relevant parameter and all parameters which are relevant match the packet, the CM or CMTS MUST forward the packet to the corresponding Service Flow (irrelevant parameters - as defined in clause C.2.1, Packet Classification Encodings - 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 CM or CMTS 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 by the CM or CMTS 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).

Cable Modem Interface Mask (CMIM): a bit mask representing the interfaces of the CM from which the CM is to classify traffic. This is a packet matching criterion in DOCSIS 3.0.

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 signalling, 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 signalling message is contained in annex C.

Attributes of QoS Classifiers include an activation state (see clause C.2.1.4.6). 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.

7.5.1.2.2 Upstream Drop Classifiers

DOCSIS 3.0 expands the concept of classifiers to encompass the filtering of upstream traffic. An Upstream Drop Classifier is a Classifier created by the CM to filter upstream traffic. If a packet matches the specified packet matching criteria of an Upstream Drop Classifier, it is then dropped.

Unlike QoS Classifiers, Upstream Drop Classifiers do not refer to a Service Flow.

The CM performs IP protocol filtering using either Upstream Drop Classifiers or IP Filters [10]. The CM reports support for Upstream Drop Classifiers in the modem capabilities encoding by sending the number of Upstream Drop Classifiers supported in the registration request (clause C.1.3.1.38). The CMTS enables Upstream Drop Classification by returning the Modem Capability Upstream Drop Classifier Support TLV with a non-zero value in the registration response. The CMTS enables IP filtering (and disables Upstream Drop Classification) by returning the Modem Capability Upstream Drop Classifier Support TLV with a value of zero in the registration response. The CMTS **MUST** allow the configuration of enabling or disabling of Upstream Drop Classification in the registration response for modems capable of using Upstream Drop Classifiers.

If Upstream Drop Classifiers are present in the configuration file, the CM **MUST NOT** include the Upstream Drop Classifier TLVs from the configuration file in the registration request message unless explicitly instructed to do otherwise via the extended MIC (see clause C.1.1.18.1.6.2).

If the CMTS enables Upstream Drop Classifiers in the registration process, the CM **MUST** filter packets using Upstream Drop Classifiers. A CM with Upstream Drop Classification enabled **MUST NOT** instantiate IP filters.

If the configuration file contains Upstream Drop Classifier Group ID(s), the CM **MUST** include the Upstream Drop Classifier Group ID(s) in the REG-REQ-MP message. If the configuration file contains Upstream Drop Classifier Group ID(s) and the registration response message contains Upstream Drop Classifiers, the CM **MUST** filter packets using the Upstream Drop Classifiers provided in the registration response message.

If the configuration file contains no Upstream Drop Classifier Group ID(s), the CM **MUST** filter packets using the Upstream Drop Classifiers provided in the configuration file. When the CM filters packets using the Upstream Drop Classifiers provided in the configuration file, the CM uses Classifier References as the Classifier IDs.

If the configuration file contains both Upstream Drop Classifier Group ID(s) and Upstream Drop Classifiers and if the registration response message contains Upstream Drop Classifiers, the CM **MUST** filter packets using the Upstream Drop Classifiers provided in the registration response message. If the configuration file contains both Upstream Drop Classifier Group ID(s) and Upstream Drop Classifiers and if the registration response message contains no Upstream Drop Classifiers, the CM **MUST** filter packets using the Upstream Drop Classifiers provided in the configuration file.

If the CMTS disables Upstream Drop Classifiers in the registration process, the CM **MUST** filter via IP filters. A CM with Upstream Drop Classification disabled **MUST NOT** instantiate Upstream Drop Classifiers. If a CM with Upstream Drop Classification disabled has received a configuration file containing both IP filters and Upstream Drop Classifiers, the CM **MUST** only instantiate IP filters. If Upstream Drop Classifiers are disabled in the registration process, the CM **MUST** reject REG-RSP or REG-RSP-MP messages or DSC-REQ messages that contain Upstream Drop Classifiers.

Like QoS Classifiers, Upstream Drop Classifiers may contain a classifier Rule Priority. The classifier Rule Priority is used for ordering the application of all Classifiers, including both Upstream Classifiers and Upstream Drop Classifiers. Explicit ordering is necessary because the patterns used by Upstream Classifiers and Upstream Drop Classifiers may overlap. The priorities need not be unique, but care must be taken within a classifier priority to prevent ambiguity in classification.

An Upstream Drop Classifier does not have associated PHS rules and is not linked to a Service Flow. The CMTS **MUST NOT** associate SF encodings or PHS Rules in an Upstream Drop Classifier in a REG-RSP, REG-RSP-MP or DSC-REQ message.

7.5.2 Object Model

The major objects of the architecture are represented by named rectangles in figure 7-13. 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 65 535 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 per logical upstream channel. A SID may be a part of a SID cluster (see clause 7.2.1.4.2.1).

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 clause 7.7). When a Service Flow is deleted, all Classifiers and any associated PHS Rules referencing it MUST also be deleted by the CM and CMTS.

The Service Class is an object that MUST 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 clause 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 clause 7.5.6.1). 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 7-13 (refer to annex L).

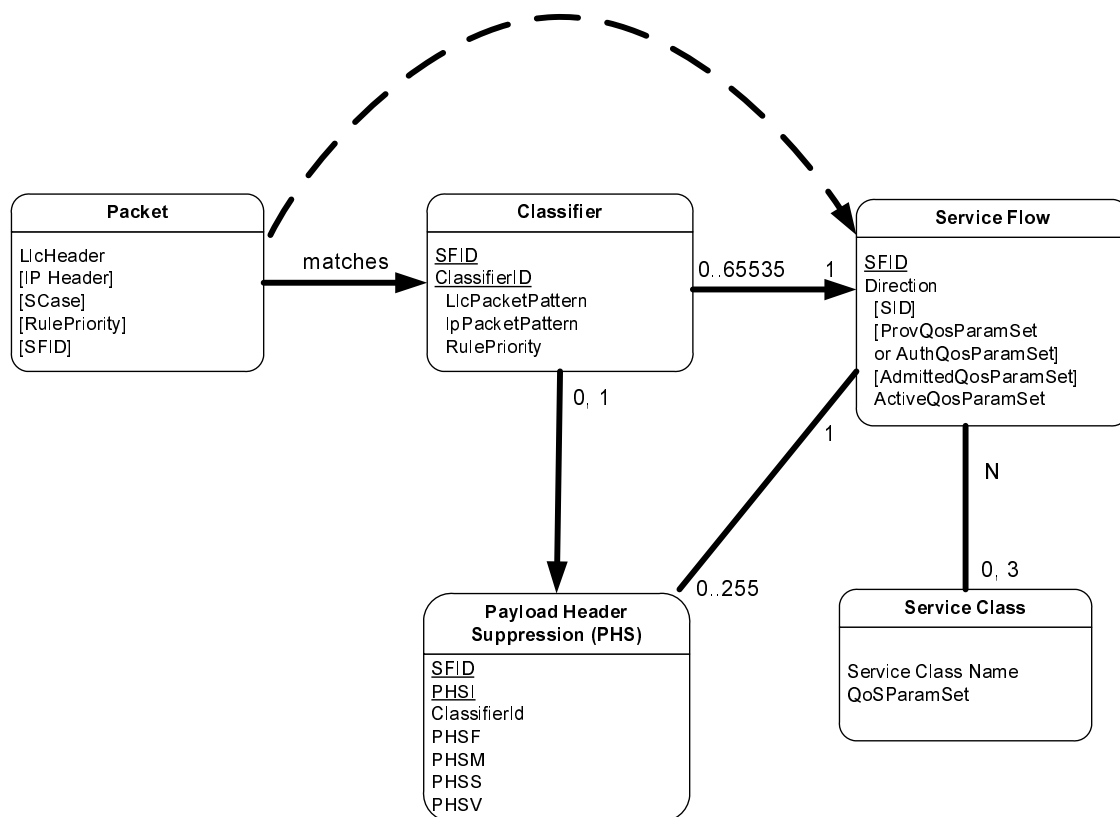


Figure 7-13: Theory of Operation Object Model

7.5.3 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 signalling 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 "unclassified" flows differently from "classified" flows with equivalent parameters.

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

- 1) By explicitly including all traffic parameters.
- 2) By indirectly referring to a set of traffic parameters by specifying a Service Class Name.
- 3) 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 from the CMTS 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.

The CMTS MAY change the QoS parameters of all downstream service flows (including both Individual and Group Service Flows) derived from a Service Class when the QoS parameters of the Service Class are changed. The CMTS MAY change the QoS parameters of all upstream service flows derived from a Service Class when those QoS parameters of the Service Class are changed. QoS parameters for downstream service flows or CMTS-enforced QoS parameters for upstream service flows, can be changed locally at the CMTS, without sending a Dynamic Service Change message to the affected CM. In order to change the CM-enforced QoS parameters of an upstream service flow, it is necessary for the CMTS to send a Dynamic Service Change message to the affected CM.

The CM-enforced QoS parameters of an upstream service flow include:

- Upstream Maximum Sustained Traffic Rate.
- Maximum Traffic Burst.
- Maximum Concatenated Burst.
- Service Flow Scheduling Type.

All other QoS parameters are CMTS-enforced.

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, REG-RSP-MP, 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.

7.5.4 Authorization

Every change to the Service Flow QoS Parameters MUST be approved by an authorization module at the CMTS. This includes every REG-REQ REG-REQ-MP 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 Service Flows in the Provisioned state. 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 may also communicate 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 may specify 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 AuthorizedQoSParamSet. Admission and activation requests from a CM that are signaled in advance by the external policy server are permitted. Admission and activation requests from a CM that are not pre-signaled 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 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.

7.5.5 States of Service Flows

It is useful to think about four states of Service Flows. This clause describes these four states of Service Flows in more detail. However, it is important to note that there are more than just these basic states.

7.5.5.1 Deferred Service Flows

A service flow may be authorized in an inactive state for subsequent admittance and activation. There are two states of deferred flows - Provisioned and Authorized.

As a result of external action beyond the scope of this specification (e.g. [25]), the CM MAY choose to authorize and/or activate a deferred 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 SID or SID Cluster(s) for an upstream Service Flow.

As a result of external action beyond the scope of this specification (e.g. [25]), the CMTS MAY choose to admit and/or activate a deferred service flow by passing the Service Flow ID as well as the SID or SID Cluster(s) and the associated QoS Parameter Sets. The CMTS MUST also provide any applicable Classifiers.

7.5.5.1.1 Provisioned Service Flows

A Service Flow may be created in the Provisioned state but not immediately activated. That is, the description of any such service flow in the TFTP configuration file contains an attribute which provisions but defers activation and admission. 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. The CMTS may deactivate the Service Flow, but SHOULD NOT delete the Service Flow during the CM registration epoch. Such a Service Flow in the Provisioned state MAY be activated and deactivated by the CMTS many times (through DSC exchanges). In all cases, the original Service Flow ID MUST be used by the CMTS when reactivating the service flow.

7.5.5.1.2 Authorized Service Flows

A Service Flow may be created in the Authorized state but not immediately activated. That is, the description of any such service flow is passed to the CMTS which authorizes but defers activation and admission (refer to clause C.2.2.3.5). The CMTS internally MUST assign a Service Flow ID for such a service flow but does not admit resources. The CMTS MAY also require an exchange with a policy module prior to admission. The CMTS MAY create, admit, activate, deactivate, de-admit and delete Service Flows which are created in the Authorized state.

7.5.5.2 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 by the CMTS (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 at the CMTS.

A Service Flow that has resources assigned to its AdmittedQoSParamSet, but whose resources are not yet completely activated, is in a transient state. A time out value MUST be enforced by the CMTS that requires Service Flow activation within this period (refer to clause 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 needs to be refreshed periodically for it to be maintained without the above timeout releasing the non-activated resources. This refresh MAY be signaled by the CMTS with a periodic DSC-REQ message with identical QoS Parameter Sets or be signaled by some internal mechanism within the CMTS outside of the scope of the present document (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".

7.5.5.3 Active Service Flows

A Service Flow that has a non-NULL set of ActiveQoSParameters is said to be in the Active state. It is requesting and being granted bandwidth for transport of data packets. A Service Flow in the Admitted state may be made active by providing an ActiveQoSParameterSet, signalling the resources actually desired at the current time. This completes the second stage of the two-phase activation model (refer to clause 7.5.5.2).

A newly created Service Flow may immediately transition to the Active state. 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 dynamically created Service Flow may be immediately transition to the Active State. In this case, two-phase activation is skipped and the Service Flow is available for immediate use upon authorization.

7.5.6 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 clause 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 clause 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 clause C.2.2.5.9).

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 or segment header off operation 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 clause C.2.1.8.3 "Ethertype/DSAP/MacType" parameter encoding and when the "type" field of the MAC Management Message Header (clause 6.4.1) matches that parameter. One exception is that the Primary SID (Multiple Transmit Channel Mode disabled) or the Ranging SID (Multiple Transmit Channel Mode enabled) MUST be used for periodic ranging, even if a classifier matches the upstream RNG-REQ message of periodic ranging. In the absence of any classifier matching a MAC Management Message, it SHOULD be transmitted by a CM or CMTS on the Primary Service Flow. Other than those MAC message types precluded from classification in clause C.2.1.8.3, 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 by the CM or CMTS. Delivery of MAC Management Messages is implicitly influenced by the attributes of the associated service flow.

7.5.6.1 Policy-Based Classification and Service Classes

As noted in annex L, 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 clause 7.7) 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 the present document. One example might be the docsDevFilterIpPolicyTable defined in the Cable Device MIB [45]. 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, classifiers that are dynamically-added by the CM or CMTS 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 the CM and CMTS SHOULD avoid the dynamic range.

7.5.7 General Operation

The CMTS MUST reject a Service Flow if the CMTS does not have the capability to support the Quality of Service parameters for the flow. For example, if the CMTS only supports certain Grant Intervals for Unsolicited Grant Service, it is required to reject a Service Flow request for a Grant Interval other than a supported value.

7.5.7.1 Static Operation

Static configuration of QoS Classifiers, Upstream Drop 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.

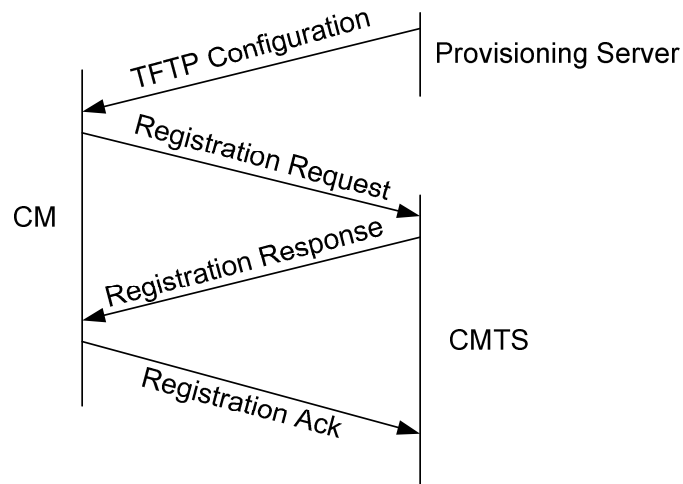


Figure 7-14: Registration Message Flow

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

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 clauses 7.5.3 and C.2.2.4.3).

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.

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 Cluster (SID when no TCC encoding is included in the Registration Response).

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.

7.5.7.2 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 7-15 and described in detail in clause 11.2.2.1.

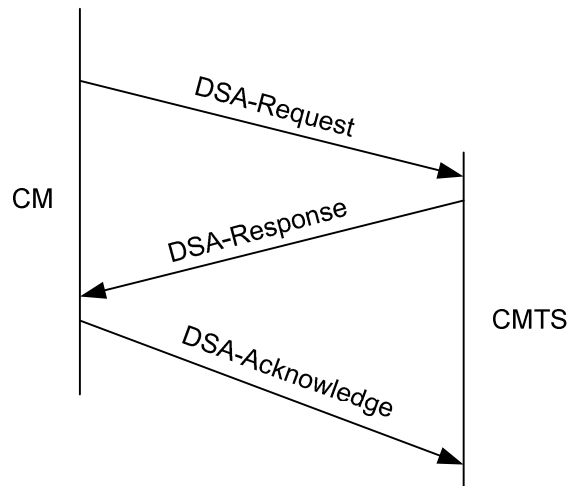


Figure 7-15: Dynamic Service Addition Message Flow - CM Initiated

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. A CM-initiated DSA-Request does not contain Upstream Drop Classifiers.

7.5.7.3 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 one or more SID Cluster Encodings, set(s) of active or admitted QoS Parameters and any required Classifier(s). A CMTS-initiated DSA-Request does not contain Upstream Drop Classifiers. The protocol is as illustrated in figure 7-16 and is described in detail in clause 11.2.2.2.




Figure 7-16: Dynamic Service Addition Message Flow - CMTS Initiated

7.5.7.4 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 clauses 11.2.3 and 11.2.4).

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 CM initiated and CMTS initiated DSC can perform the following actions:

- Add, replace or delete QoS classifiers.
- Create, add parameters to or delete PHSI-indexed PHS rules.

- Create or delete DSID-indexed PHS rules.

The CMTS-initiated DSC can also add, replace or delete Upstream Drop Classifiers. The CMTS **MUST** reject a CM-initiated DSC containing a DSC action to add, replace or delete an Upstream Drop Classifier. The DSC cannot be used to change Service Flow SID Clusters. The CM **MUST** reject a CMTS-initiated DSC which attempts to change Service Flow SID Clusters.

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 clause C.2.2.3.5) 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 clause 7.5.1.1). 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.

The DSD cannot be used to delete Upstream Drop Classifiers.

7.5.8 QoS Support for Joined IP Multicast Traffic

This clause describes a standard configuration and implementation of QoS for downstream IP multicast traffic that is joined dynamically by a multicast host or statically joined via CMTS configuration. The mechanism for providing QoS to a group of CMs is similar to the mechanism for providing it to an individual CM: the highest priority classifier that matches a downstream packet identifies the service flow for scheduling the packet. In the case of multicast traffic, the classifiers are called "Group Classifier Rules" (GCRs) and the service flows are called Group Service Flows (GSFs). GCRs and GSFs are associated with a Downstream Channel Set (DCS), which is either a single downstream channel or a downstream bonding group of multiple downstream channels. A MAC Domain is considered to have Individual Classifier Rules and Individual Service Flows associated with an individual Cable Modem as well as Group Classifier Rules (GCRs) and Group Service Flows (GSFs) associated with a Downstream Channel Set (DCS). GCRs and GSFs have the same attributes and are described in the same MIB tables as Individual Classifier Rules and Individual Service Flows.

This clause describes QoS only for joined IP multicast sessions. This includes dynamically joined sessions using multicast management protocol such as IGMP/MLD as well as statically joined sessions using Static Multicast Session Encodings in REG-REQ(-MP) (see clause C.1.1.27). The mechanism by which the CMTS provides QoS for other downstream broadcast and layer 2 multicast traffic is CMTS vendor specific, although certain CMTS requirements for this traffic are described below.

7.5.8.1 Overview

An object model diagram that describes multicast QoS operation is depicted in figure 7-17.

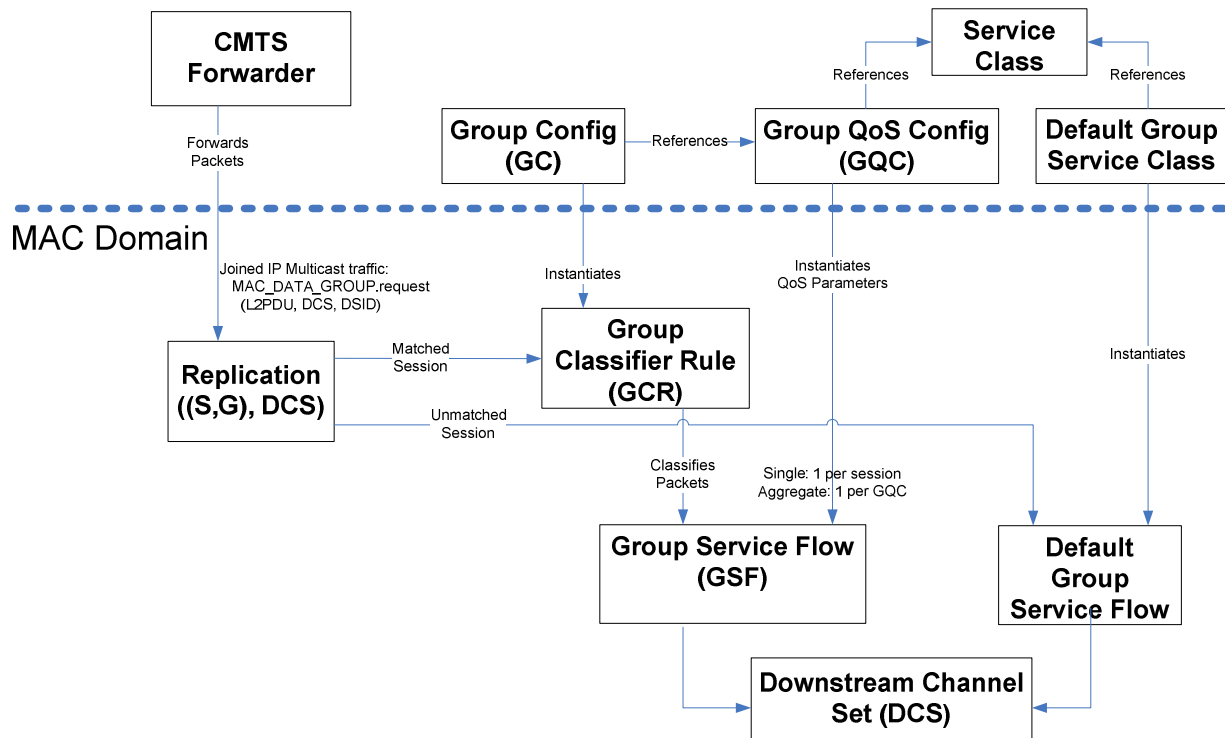


Figure 7-17: IP Multicast QoS Object Model Diagram

The operational model of the CMTS as described in annex L is that a CMTS Forwarder submits a `MAC_DATA_GROUP.request` primitive to a MAC Domain in order to replicate a downstream multicast frame on a Downstream Channel Set of the MAC Domain:

- `MAC_DATA_GROUP.request(MAC frame, DCSid, DSID).`

Where:

- MAC frame consists of the layer 2 Ethernet MAC packet from Destination Address through CRC.
- DCSid is an index of a Downstream Channel Set that corresponds to either a single downstream channel or a downstream bonding group of multiple channels.
- DSID is a Downstream Service Identifier that identifies the group of Cable Modems to which the CMTS Forwarder is transmitting the packet.

Because the CMTS Forwarder supplies a MAC frame to the MAC Domain, it is considered to have mapped the IP destination group address into the standardized layer 2 Ethernet group MAC address for that IP destination group.

As depicted in figure 7-17, each DCS is considered to have one or more Group Classifier Rule (GCR) objects associated with it. GCRs are considered to be classifiers of the MAC Domain with the same attributes as the classifiers defined for individual cable modems. Different DCSs have different sets of GCRs.

The DSID is intended to identify the combination of a specific IP multicast session and DCS. The CMTS assigns a different DSID to the same multicast session replicated on different DCSs. The CMTS assigns a different DSID to each different multicast session replicated to the same DCS. A DSID value is unique per MAC Domain (i.e. the CMTS can re-use the same DSID value for two different multicast sessions being transmitted on two different MAC domains).

When the CMTS Forwarder requests a MAC Domain to transmit a joined IP multicast session packet on a particular DCS, the MAC Domain compares the packet against the list of Group Classifier Rules (GCRs) associated with the DCS of the request.

Each GCR refers to a single Group Service Flow (GSF) instantiated on the DCS. A Group Service Flow is a downstream Service flow with the same QoS Parameter Sets as an Individual Service Flow (ISF) created for an individual cable modem. A GSF is always active; its Provisioned, Admitted and Active QoS Parameter Sets are the same set. GSFs are not communicated to CMs. A GSF is intended to be assigned to the same DCS for the duration of its lifetime. A GSF is not considered to be autonomously load balanced to other DCSs. When the CMTS changes the replication of a particular IP multicast session to a different DCS, the session is considered to be scheduled on a different GSF for the new DCS.

GCRs, like individual classifier rules, have a rule priority. If the multicast packet matches more than one GCR then the CMTS uses the GCR with highest rule priority to select the GSF for transmitting the packet. If more than one matching GCRs have the same highest rule priority, the GCR used by the CMTS is vendor-specific.

If the packet does not match any GCR, the CMTS forwards it to a Default Group Service Flow that is instantiated with QoS parameters from the identified Default Group Service Class for the CMTS.

The Group Service Flow identified for a downstream packet controls the QoS and provides the statistics for accounting for the transmission of the packet in the MAC Domain.

7.5.8.2 Group Configuration and Group QoS Configuration Tables

For IP Multicast QoS, a cable operator controls the creation of GCRs and GSFs by configuring entries in Group Configuration (GC) and Group QoS Configuration (GQC) tables. These tables only configure the QoS for IP Multicast sessions; they do not control how CMTS replicates IP Multicast Sessions on DCSs. Replication of IP Multicast sessions is determined based on joiners to IP Multicast Sessions. Configuration of how the CMTS replicates to DCSs (e.g. whether the CMTS replicates certain sessions to downstream bonding groups or to a single downstream channel) is vendor-specific.

The object model representation of the Multicast QoS configuration and the associated reporting objects are defined in Multicast QoS Configuration clause of the Multicast Requirements annex in [10]. Several new objects have been defined to allow the Operator to configure and determine the replication of each Multicast group that is forwarded on a given DCS. These objects include:

- **Group Configuration:** An object that defines the matching criteria for multicast sessions that have been configured for specific QoS treatment. This object is used to define the Group Classifier Rules (GCR) that will place traffic on a given Group Service Flow (GSF). The object also defines if PHS and encryption are needed for a given multicast session.
- **Group QoS Config:** An object to assign the specific QoS attributes of a Group Service Flow (GSF) that uses Service Class Names to define the specific QoS treatment that a given multicast session requires.
- **Group PHS Config:** An object for assigning the PHS rules associated with a multicast session.
- **Group Encryption Config:** An object for defining the rules for encrypting multicast sessions. This table does not control the encryption of sessions required for Isolation by the CMTS, only the need for a given multicast session to be encrypted regardless of isolation.
- **Replication Session:** An object used to report the status and forwarding of all multicast sessions actively being forwarded on all DCS in a CMTS.

Operators can configure the rules for QoS, PHS and Encryption by creating entries in the Group Configuration Object. QoS is configured using the GC and GQC objects. Encryption for specific multicast sessions is configured using the GC and GroupEncryption Objects. PHS rules for Multicast sessions are defined using the GC and GroupPHS objects.

The following tables give some examples of the Session Range and Differentiated Services (ToS) specifiers in the Group Configuration object.

Table 7-8: Examples of Group Configuration Session Ranges

(*,G)	All IP multicast sessions to a specific group address (G)
(*,G range)	All sessions to a range of groups (Gs)
(S,*)	All sessions from a specific source host (S)
(S,G)	A specific session from source (S) to group (G), i.e. a Source Specific Multicast (SSM) session
(S range, G)	All sessions from a masked range of sources (Ss) AND to a particular group (G)
(S range, G range)	All sessions from a range of sources (Ss) to a range of groups (Gs)
(*,*)	All IP multicast sessions

Table 7-9: Examples of IP DS Field Ranges

(*,*,*)	All IPv4 TOS values or all IPv6 Traffic Class values with a mask of all bits. This will match all packets within the session range defined in the GC entry.
(L,L,*)	The single IPv4 TOS values or IPv6 Traffic Class equal to L with all bits in the corresponding field being valid. This will match only those packets in the session range defined in the GC entry whose DS field exactly matches "L".
(L,H,*)	All IPv4 TOS values or all IPv6 Traffic Class values within the range of L to H with all bits in the corresponding field being valid. This will match only those packets in the session range defined in the GC entry whose DS field is greater than or equal to "L" and less than or equal to "H".
(L,L, B)	The single IPv4 TOS values or IPv6 Traffic Class equal to L only considering the bits of the corresponding field denoted by the bits in the mask represented by B. This will match only those packets in the session range defined in the GC entry whose DS field logically AND-ed with bit mask B exactly matches "L".
(L,H, B)	All IPv4 TOS values or all IPv6 Traffic Class values within the range of L to H only considering the bits of the corresponding field denoted by the bits in the mask represented by B. This will match only those packets in the session range defined in the GC entry whose DS field logically AND-ed with bit mask B is greater than or equal to "L" and less than or equal to "H".

7.5.8.3 Instantiating Group Classifier Rules and Group Service Flows

The CMTS MUST take the following steps to instantiate GCRs and GSFs for controlling the QoS of dynamically or statically joined IP Multicast sessions:

- 1) When a client within a MAC Domain joins a multicast session, the CMTS determines to which Downstream Channel Set (DCS) it will replicate the packets of that session in order to reach the multicast client. The DCS may be a single downstream channel or a downstream bonding group of multiple downstream channels. The CMTS assigns a Downstream Service ID (DSID) for the replication of a particular session onto a particular DCS. The CMTS SHOULD select a DCS for replication such that the Service Class named in the GQC entry referred to by the GC entry matched by the session has a Required Attribute Mask with attributes that are also set for the DCS and a Forbidden Attribute Mask with attributes that are not set for the DCS. The attributes for a DCS are either configured directly (for an individual channel or a provisioned downstream bonding group) or derived from the component channels of the DCS and the Attribute Aggregation Mask for the Service Class (for a dynamically created bonding group).
- 2) The CMTS determines the set of GC entries whose Session Range matches the new (S, G) session. If more than one GC entry matches, the CMTS selects the GC entry with the highest Rule Priority. If more than one matching GC entry has the same highest Rule Priority, then all the GC entries matching the highest Rule Priority are selected for instantiating GCRs and GSFs. If no GC entry has a Session Range that matches the new session, the CMTS does not create any new Group Classifier Rule (GCR) or Group Service Flow (GSF) for the session; in this case the packets of the session are transmitted using the default GSF for the DCS, as described below. The CMTS creates GCR and GSF entries only when there is a reference to a valid (non zero) GQC entry from the GC entry which matches a session. However, irrespective of whether GCR and GSF entries are created for the matched session or not, the encryption and PHS rules are applied to the session if there are references in the matched GC entry to valid encryption and PHS rules.

- 3) When a matching GC entry is selected for the first joiner of a session on the DCS, the CMTS may instantiate a Group Classifier Rule (GCR) for classifying the session's packets, based on whether the QoS Control of the selected GC entry is "Single-Session" or "Aggregate-Session":
- For a QoS Control of "Single-Session", the CMTS always creates a GCR with the specific Group (G) destination IP address as a criterion. If the session was joined with a protocol supporting Source Specific Multicast (SSM), the GCR also contains the particular Source (S) IP address. A single-session GQC entry thus creates a single-session GCR. When other unique (S, G) sessions are joined that match the session range of the single-session GC entry, the CMTS creates a separate GCR for each session. The CMTS creates only a single-session GCR and GSF for all CMs with joiners reached by the same session on the same DCS.
 - For a QoS Control of "Aggregate-Session", the CMTS creates a GCR with the same session range (e.g. S-range, G-range) criteria as the GC entry itself, if such a GCR has not already been created on a DCS. The GCR created by an "Aggregate-Session" GC entry classifies an aggregate of multiple multicast sessions. The CMTS creates at most one Aggregate-Session GCR and GSF on a DCS for each Aggregate-Session GC entry.

In both cases, the instantiated GCR uses the same Rule Priority as specified for the Rule Priority of the selected GC entry itself.

The CMTS may implement vendor-specific configuration that controls the mapping of Source Specific Multicast (SSM) network sessions to multicast joins performed with an Any Source Multicast (ASM) protocol (i.e. that requests joining a session identified only by its destination group address). This vendor-specific configuration can also determine which GC entries with explicit source addresses apply to ASM joins.

- 4) The CMTS then may create a Group Service Flow (GSF) for the new session replication. The Service Class named in a GQC entry provides the template for the QoS parameters assigned to the GSF. A valid GQC entry references an existing Service Class Name in the CMTS Service Class Table. Typical QoS parameters for a GSF include Minimum Reserved Traffic Rate and the Maximum Sustained Traffic Rate. A Group Service Flow is assigned to a single DCS and remains assigned to that DCS for the duration of its instantiation. If the attribute mask for a DCS does not match all required attributes or does match any forbidden attribute of the Service Class of the GSF, the CMTS MUST log an event and update the MIB to report an "attribute assignment failure" event when it creates the GSF. If the individual channel Service Flow Attribute Masks or the Aggregate Service Flow Attribute Masks are changed and these changes conflict with the Service Flow Attribute Masks defined in the SCN, the CMTS MUST note the error in the event log. In this case, the CMTS may need to move the replication to in order to satisfy the defined Service Flow Attribute Masks defined in the SCN. The QoS Control of the GQC entry determines how the CMTS instantiates GSFs for the GQC entry:
- For a QoS Control of "Single-Session", the CMTS creates a GSF on a DCS for each single session, that is, each unique combination of source IP address S (for an SSM session) and destination group IP address G.
 - When a single GC entry that matches a range of multicast sessions references a GQC entry with a QoS Control of "Aggregate-Session", the CMTS creates a GSF on a DCS for the first multicast session matching that GC entry. When another session matches the same Aggregate-Session GC entry, the CMTS does not create another GSF and does not create another GCR for the existing GSF. In this case, the CMTS associates a single GSF and a single GCR for all multicast sessions matching an Aggregate-Session GC entry. Thus, all the multicast sessions that match a GC entry (e.g. S-range, G-range) share the same bandwidth allocated for the GSF, instead of creating a separate GSF for each multicast session that matches a GQC entry.
 - When multiple GC entries refer to the same GQC entry with QoS Control of "Aggregate-Session", the CMTS creates only one GSF. For the first multicast session matching a GC entry, the CMTS creates a GSF and a GCR corresponding to the matched GC entry. For subsequent multicast session matching another GC entry that references the same GQC entry, the CMTS creates a new GCR entry and associates the GCR entry with the existing GSF.

- 5) The CMTS will maintain GCRs and GSFs on a DCS for as long as it replicates multicast sessions that use them. The CMTS may discontinue replication of a session onto a DCS either because the last joiner has left or because it elects to replicate the session to a different DCS. When the CMTS discontinues forwarding of a multicast session to a DCS, it deletes any Single-Session GCR and single-session GSF it had created for the multicast session. When the CMTS discontinues forwarding of the last of the multicast sessions for which it had created an Aggregate-Session GCR, the CMTS deletes the Aggregate-Session GCR. When the CMTS deletes the last GCR that refers to an Aggregate-Session GSF, it deletes the aggregate-session GSF itself.
- 6) A CMTS may create GCRs and GSFs for IP Multicast sessions in a vendor-specific manner. The CMTS will assign the rule priority attribute of a vendor-specific GCR to be in the range 64 to 191. This permits GCRs instantiated from the operator-configured GC entry to have either a lower priority (0 to 63) or higher priority (192 to 255) than the vendor-specific GCR entries.

Cable operators need to take great care when assigning the bandwidth attributes of Group Service Flows for aggregate sessions to avoid service flows that do not provide enough or reserve too much bandwidth for the aggregate sessions. When the bandwidth of each multicast session to be aggregated is known, the cable operator can configure `AggregateSessionLimit` to control the maximum bandwidth of the aggregate GSF. When the bandwidth of each multicast session to be aggregated is not known, the cable operator can configure the downstream maximum sustained traffic rate (clause C.2.2.5.2) of the aggregate GSF.

When a client joins an IP Multicast Session, there may be insufficient resources to schedule traffic from the session on a GSF (Single-Session or Aggregate-Session). The CMTS behavior in this case is vendor-specific.

CMTS operation concerning invalid GC and GQC entries is vendor-specific. The CMTS may prevent the creation of an invalid GC or GQC entry, e.g. one that contains a name for a Service Class that does not exist. The CMTS may prevent the deletion of configuration objects that would result in "dangling references", e.g. the deletion of a Service Class referenced by a GQC entry.

7.5.8.3.1 Examples of GCR and GSF Instantiation

EXAMPLE 1:

This first example covers classifying multiple multicast sessions matching two different GC entries into a single shared GSF (two GC entries referencing a single GQC entry of type "Aggregate-Session").

In this example, a stockbroker "Broker A" has contracted with the cable operator to provide pushed multicast stock quotes. Each stock symbol issue has a separate IP multicast destination group and there are potentially hundreds of such groups. The broker has identified two IP multicast source hosts S1 and S2 that generate these stock quotes. The cable operator has agreed to provide at least 20 Kbps of bandwidth but no more than 100 Kbps for the aggregate of 10 multicast sessions.

The operator configures two GC entries. Each GC entry applies to all MAC Domains and to all Downstream Channel Sets of those MAC Domains. Entry GC1 contains the IP multicast session range (S1,*), IP DS Range (0x00,0xFF,0xFF) to match all markings. Entry GC2 contains the session range (S2,*), IP DS Range (0x00,0xFF,0xFF) to match all markings. Both GC1 and GC2 refer to a GQC1 entry with QoS Control "Aggregate-Session", `AggregateSessionLimit` of 10 and Service Class named "BrokerMcast".

- Group Config Entries:
 - GC1: Session Range=(S1,*), IP DS Range=(0x00,0xFF,0xFF), GroupQoSConfigId=GQC1.
 - GC2: Session Range=(S2,*), IP DS Range=(0x00,0xFF,0xFF), GroupQoSConfigId=GQC1.
- Group QOS Config Entry:
 - GQC1: QoS Control=Aggregate-Session, AggregateSessionLimit =10, SCN="BrokerMcast".

The operator configures the Service Class Table with a class named "BrokerMcast" with a Minimum Reserved Traffic Rate of 20 Kbps and a Maximum Sustained Traffic Rate of 100 Kbps.

- ServiceClassTable Entry:
 - BrokerMcast: MinReserved = 20 000 bps, MaxSustained = 100 000 bps.

When the first joiner of any multicast session from S1, say (S1, G1), joins on a particular MAC domain, the CMTS selects the Downstream Channel Set to reach that joiner, creates GSF1 on that DCS and has GCR1 that references GSF1. GCR1 has the same (S1,*) classification criteria as GC1:

- GCR1: (S1,*) → GSF1.

When a joiner to a second multicast session from S1, say (S1, G2), joins on the MAC domain and the CMTS elects to distribute the session to the same DCS, the CMTS does not create any new GCR-it keeps GCR1-and it does not create any new GSF-it keeps GSF1. This is because GC1 references a GQC entry of type "Aggregate-Session".

When the first joiner for any multicast session from S2, say (S2, G20), joins on the MAC domain and the CMTS elects to distribute the session to the same DCS, the CMTS does not create a new GSF, but it does create a new GCR2 that references the same GSF1 it created earlier. This is because the GC1 and GC2 both reference the same GQC entry, GQC 1. GCR2 also uses the same wild-card criteria as GC2:

- GCR2: (S2,*) → GSF1.

The MAC Domain has two GCRs-GCR1 and GCR2-that each reference the same GSF-GSF1.

Since GQC 1 entry specified AggregateSessionLimit of 10, only 10 multicast sessions matching GCR1 and GCR2 can be transmitted simultaneously using GSF1.

EXAMPLE 2:

This second example covers classifying multiple multicast sessions matching two different GC entries into two separate shared GSFs (two GC entries referencing two different GQC entries of type "Aggregate-Session").

The cable operator from Example 1 contracts with two additional stockbrokers "Broker B" and "Broker C" for the same IP multicast push service. Broker B has a single IP multicast source S3 and Broker C has a single IP multicast source S4. Each broker is promised the same QoS service level agreement, namely at least 20 kbps and at most 100 kbps for the aggregate of 10 joined multicast sessions for each broker.

The operator configures two GC entries corresponding to each broker's IP multicast sources S3 and S4 with IP DS Range of (0x00,0xFF,0xFF) to match all IP class markings. Each GC entry references a separate GQC entry with QoS Control of "Aggregate-Session", because a separate shared GSF needs to be created for each GC entry. This allows each GSF to meet the QoS service level agreement with the individual broker. Both the GQC entries have AggregateSessionLimit of 10 and reference the same Service Class as Example 1:

- Group Config Table Entries:
 - GC3: Session Range=(S3, *), IP DS Range (0x00,0xFF,0xFF), GroupQoSConfigId=GQC2.
 - GC4: Session Range=(S4, *), IP DS Range (0x00,0xFF,0xFF), GroupQoSConfigId=GQC3.
- Group QoS Config Table Entries:
 - GQC2: QoS Control=Aggregate-Session, AggregateSessionLimit =10, SC="BrokerMcast".
 - GQC 3: QoS Control=Aggregate-Session, AggregateSessionLimit =10 SC="BrokerMcast".

When the first joiner joins any session from S3, for example (S3,G3), the CMTS creates GCR3 using the same wild-card range criteria as GC3. The CMTS also creates a new GSF2 for the aggregate set of 10 multicast sessions from S3 and has GCR3 point to GSF2:

- GCR3: (S3,*) → GSF2.

When the first joiner joins any session from S4, for example (S4, G4), the CMTS creates GCR4 with the same wild-card range criteria of GC4 and has it reference a newly created GSF3 for the aggregate of 10 multicast sessions from S4:

- GCR4: (S4,*) → GSF3.

In this case, the QoS received by Broker B's multicast sessions (from S3) is independent of the QoS received by Broker C's sessions (from S4) because they each have a separate GSF on the Downstream Channel Set. This is required as the cable operator needs to honor the QoS service level agreement established with each broker.

Also note that since both GQC 2 and GQC 3 specified AggregateSessionLimit of 10, only 10 multicast sessions matching GCR3 can be simultaneously transmitted using GSF2 and only 10 multicast sessions matching GCR4 can be simultaneously transmitted using GSF3.

All of the GCRs-GCR1, GCR2, GCR3 and GCR4-are "Aggregate-Session" GCRs because their classifier criteria matches a range of multiple (S, G) IP sessions. All of the GSFs-GSF1, GSF2, GSF3-are "Aggregate-Session" GSFs because they transmit multiple IP multicast sessions, using three separate, shared GSFs. If any IP multicast session that is being transmitted on a shared GSF, sends excessive traffic, all of the IP multicast sessions sharing that particular GSF can be affected. In this case, however, the QoS received by the IP multicast sessions aggregated on other shared GSFs would not be affected.

EXAMPLE 3:

This third example covers creating a separate, dedicated GSF for individual IP Multicast session: (a GC entry referencing a Single Session GQC Entry).

In this example, each IP multicast session represents a switched broadcast IP Video transmission, e.g. a standard definition MPEG-2 stream of approximately 3,75 Mbps originated by a Cable Operator-provided IP Video server S6. Once an IP video stream has been assigned to a particular Downstream Channel Set, it must not be affected by any other unicast or multicast traffic. This is a requirement for "Single-Session" QoS, where each individual session has its own GCR and GSF.

The operator configures the IP DS Range to (0x00,0xFF,0xFF) to match all class markings, since IP Class markings are not required for this service definition. The GC entry references a GQC entry with QoS Control of "Single-Session" and Service Class named "Mpeg2SD".

- Group Config Table Entry:
 - GC5: Session Range=(S6, *), IP DS Range (0x00,0xFF,0xFF), GroupQoSConfigId=GQC4.
- Group QoS Config Table Entry:
 - GQC 4: QoS Control="Single-Session", SC="Mpeg2SD".

The operator configures the Service Class Table with a class named "Mpeg2SD" with both the Minimum Reserved Rate and Max Sustained Rate to be 4 Mbps and Max Burst size to be 1 000 000 bytes (1 Mbyte).

- The Service Class Table Entry:
 - Mpeg2SD: MinReserved = 4 000 000 bps, MaxSustained = 4 000 000 bps, MaxBurst = 1 000 000 bytes.

When the first host joins an individual session matching (S6,*), for example (S6, G5), the CMTS creates a new Group Classifier Rule GCR5 and a new Group Service Flow GSF4 on the Downstream Channel Set. The GCR matches the single session of the GC entry, namely (S6, G5):

- GCR5: (S6, G5) -> GSF4.

When a host joins a second session from S6, for example (S6, G6), the CMTS creates a new Single-Session GCR6 and it creates a new GSF5 for the session because the GC entry references a GQC entry of type "Single-Session":

- GCR6: (S6, G6) -> GSF5.

NOTE 1: Two important differences between this example and from the two "Aggregate-Session" examples.

NOTE 2: Each instantiated GCR has criteria that matched the particular, specific session joined.

NOTE 3: Each instantiated GCR references a newly-created GSF for the particular, specific session.

EXAMPLE 4:

This fourth example covers classifying multiple multicast sessions with the same Session Range but different IP DS Ranges into two separate shared GSFs (two GC entries with same Session range but different IP DS Ranges referencing two different Aggregate-Session GQC entries).

Similar to Example 1 this example has, a stockbroker "Broker A" who contracted with the cable operator to provide pushed multicast stock quotes and multicast IPTV NEWS feeds. Each stock issue has a separate IP multicast destination group and there are potentially hundreds of such groups. Each IPTV NEWS feed also has a separate destination group and there are potentially tens of such groups. The source of both the stock quote sessions and the IPTV NEWS sessions is not known but the server will mark the IPTV service and stock quote sessions with a different IPv4 TOS values to distinguish them. The stock quote service will be marked with an IPv4 TOS value of 1 and the IPTV NEWS feed will be marked with a IPv4 TOS value of 2. All other IPv4 TOS values are not expected but the operator has agreed to put those flows into their default class of service. The cable operator has agreed to provide at least 20 kbps of bandwidth on each downstream channel for the aggregate of all stock quote related multicast sessions, but no more than 100 Kbps for the aggregate of all stock quote related sessions. The cable operator has agreed to provide at least 4 Mbps of bandwidth on each downstream channel for the aggregate of all IPTV NEWS related sessions, but no more than 10 Mbps for the aggregate of all IPTV NEWS related sessions.

The operator configures two GC entries. Each GC entry applies to all MAC Domains and to all Downstream Channel Sets of those MAC Domains and applies to all IP multicast sessions (*,*) as the group and source are unknown to the operator. Entry GC6 contains IP DS Range (1,1,0xFF) and references a GQC entry with, QoS Control of "Aggregate-Session" and Service Class of "Broker Stock". Entry GC7 contains IP DS Range (2,2,0xFF) to map all other IPv4 TOS values and references a GQC entry with QoS Control "Aggregate-Session" and a different Service Class - "Broker IPTV".

- Group Config Table Entries:
 - GC6: Session Range=(*, *), IP DS Range= (1,1,0xFF), GroupQosConfigId=GQC5.
 - GC7: Session Range=(*, *), IP DS Range = (2,2,0xFF), GroupQosConfigId=GQC6.
- Group QoS Config Table Entries:
 - GQC5: QoS Control=Aggregate-Session, SC="BrokerStock.
 - GQC6: QoS Control=Aggregate-Session, SC="BrokerIPTV".

The operator configures two new Service Classes in the Service Class Table. The first is "BrokerStock" with a Minimum Reserved Traffic Rate of 20 Kbps and a Maximum Sustained Traffic Rate of 100 Kbps. The second is "BrokerIPTV" with a Minimum Reserved Traffic Rate of 4 Mbps and a Maximum Sustained Traffic Rate of 10 Mbps.

- ServiceClassTable:
 - BrokerStock: MinReserved = 20 000 bps, MaxSustained = 100 000 bps.
 - BrokerIPTV: MinReserved = 4 000 000 bps, MaxSustained = 10 000 000 bps.

When the first joiner of any session from any source say S7, joins any group on a particular MAC domain, the CMTS selects the Downstream Channel Set to reach that joiner, creates GSF6 and GSF7 on that DCS and has GCR7 reference GSF6 and GCR8 reference GSF7. GCR7 and GCR8 have the same (*,*) criteria as GC6 but different IP DS criteria:

- GCR7: (*,*) IP DS (1,1,0xFF) → GSF6.
- GCR8: (*,*) IP DS (2,2,0xFF) → GSF7.

When a joiner to a second session from any source joins any group on the MAC domain and the CMTS elects to distribute the session to the same DCS, the CMTS does not create any new GCR--it keeps GCR7 and GCR8 --and it does not create any new GSF--it keeps GSF6 and GSF7. This is because GC6 and GC7 reference GQC entries of type "Aggregate-Session".

IP Multicast packets forwarded by the CTMS to any MAC domain for any session will use GCR7 if the IP DS field is set to 1 and hence will be transmitted using GSF6. IP Multicast packets forwarded by the CMTS to any MAC domain for any session with IP DS field = 2 will instead use GCR8 and hence be transmitted using GSF7. IP Multicast packets that do not contain an IP DS field of 1 or 2 will be forwarded using the default GSF since no GCR is defined for those IP DS field values.

NOTE 4: Since no AggregateSessionLimit is specified for GQC entries in this example, there is no limit on how many multicast sessions are transmitted simultaneously using GSF6 and GSF7.

An important differences between this example and the previous examples is that this example shows multiple GC entries with the same Session Range (*,*) but different IP DS Ranges. This tells the CMTS to create multiple GCRs for a single join, one for each different IP DS range. Because only IP DS 1 and 2 were configured in the GQC table and all other IP DS values will go to the default GSF.

EXAMPLE 5:

This fifth example covers classifying multicast sessions matching two GC entries with the same Session range but different IP DS Range into separate Single Session GSFs (two GC entries with same Session range but different IP DS Ranges referencing two different Single-Session GQC entries).

Similar to example 3 in this example, each IP multicast session represents a switched broadcast IP Video transmission, e.g. a standard definition MPEG-2 stream of approximately 3,75 Mbps or a high definition MPEG-2 stream of approximately 8 Mbps originated by a Cable Operator-provided IP Video server S8. Once an IP video stream has been assigned to a particular Downstream Channel Set, it must not be affected by any other unicast or multicast traffic. This is a requirement for "single-session" QoS, where each individual session has its own GCR and GSF. However there is one twist to this example, the Cable Operator wants high definition TV, labeled by the server with an IP TOS value of 255, to be guaranteed higher bandwidth than standard definition TV as it requires more bandwidth for the higher quality. The Cable Operator has identified one IP multicast source host S8 for bother standard definition and high definition TV streams.

The operator configures two GC entries. Each GC entry applies to all MAC Domains and to all Downstream Channel Sets of those MAC Domains. Entry GC8 contains the IP multicast session range (S8,*), with IP DS Range (0,254,255) and references GQC 7 with QoS Control "Single-Session", Service Class "Mpeg2SD". Entry GC9 contains the same session range (S8, *), but contains IP DS Range of (255,255,255) to recognize the High definition TV flows and references GQC 8 with QoS Control "Single-Session", Service Class = " Mpeg2HD".

- Group Config Table Entries:
 - GC8: Session Range=(S8, *) IP DS Range=(0,254,255), GroupQoSConfigId=GQC7.
 - GC9: Session Range=(S8, *) IP DS Range=(255,255,255), GroupQoSConfigId=GQC8.
- Group Qos Config Table Entries:
 - GQC7: QoS Control="Single-Session", SC=" Mpeg2SD".
 - GQC8: QoS Control="Single-Session", SC=" Mpeg2HD".

The operator uses the "Mpeg2SD" Service Class defined in Example 3 above and configures a new Service Classes in the Service Class Table for the High definition TV streams. The new Service Class is "Mpeg2HD" and has a Minimum Reserved Traffic Rate of 8 Mbps, a Maximum Sustained Traffic Rate of 16 Mbps and a Maximum Burst Size of 1 MByte.

- ServiceClassTable:
 - Mpeg2HD: MinReserved = 8 000 000 bps, MaxSustained = 16 000 000 bps, MaxBurst = 1 000 000 bytes.

When the first joiner of any session from S8, say (S8, G9), joins on a particular MAC domain, the CMTS selects the Downstream Channel Set to reach that joiner, creates GSF8 and GSF9 on that DCS and has GCR9 reference GSF8 and GCR10 reference GSF9. GCR9 and GCR10 have the same specific (S8,G9) criteria but different IP DS criteria, since GC8 and GC9 referenced GQC entries of type "Single-Session":

- GCR9: (S8,G9) IP DS (0,254,255) → GSF8.
- GCR10: (S8,G9) IP DS (255,255,255) → GSF9.

When a joiner to a second session from S8, say (S8,G10), joins on the MAC domain and the CMTS elects to distribute the session to the same DCS, the CMTS creates a new set of GCRs and GSFs for the new group (G10). This is because GC8 and GC9 referenced GQC entries of type "Single-Session":

- GCR11: (S8,G10) IP DS (0,254,255) → GSF10.
- GCR12: (S8,G10) IP DS (255,255,255) → GSF11.

IP Multicast packets forwarded by the CTMS to any MAC domain for session (S8,G10) will use GCR11 if the IP DS field is set to any value other than 255 and hence will be transmitted using GSF10. IP Multicast packets forwarded by the CMTS to any MAC domain for session the same session (S8,G10) but with IP DS field = 255 will instead use GCR12 and hence be transmitted using GSF11.

NOTE 5: Like Example 4, this example has multiple GCs with the same Session Range but different IP DS Ranges causing a single join to create multiple GCRs, one for each IP DS Range.

NOTE 6: The two important differences between this example and the Aggregate-Session example #4:

- Each instantiated GCR has criteria that matched the particular, specific session joined.
- Each instantiated GCR references a newly-created separate, single-session GSF for the particular, specific session.

EXAMPLE 6:

This sixth example covers classifying multicast sessions matching one Single Session GC entry with a specific IP DS field range into a separate Single-Session GSFs leaving the other remainder multicast sessions to use the default GSF (single GC entry with a specific IP DS field referencing a "Single-Session" GSF).

In this example each multicast session represents an IPTV feed. The operator has configured their local content servers to use IPv6 Traffic Class = 6. Each stream from their local server is approximately 3,75 Mbps, but they come from various servers so the source is unknown. Once an IP video stream has been assigned to a particular Downstream Channel Set, it must not be affected by any other unicast or multicast traffic. This is a requirement for "single-session" QOS, where each individual session has its own GCR and GSF. The operator also wishes to treat other multicast traffic as best effort without guarantee. The operator has configured their network such that other multicast traffic will never arrive at the CMTS with an IPv6 Traffic Class = 6.

The operator configures one GC entry for its local IPTV sessions. The GC entry applies to all MAC Domains, to all Downstream Channel Sets of those MAC Domains and to all IP multicast sessions, as the group and source are unknown to the operator. Entry GC10 contains IP DS Range (6,6,255) and references the GQC9 with QoS Control "Single-Session" and Service Class "Mpeg2SD".

- Group Config Table Entries:
 - GC10: Session Range=(*, *), IP DS Range = (6,6,255), GroupQosConfigId=GQC9.
- GroupQosTable:
 - GQC9: QoS Control="Single-Session", SC="Mpeg2SD".

By creating no other GC entries the operator configures the CMTS to use a default, best effort GSF for all other IP DS field values.

The operator uses the service class "Mpeg2SD" defined in example 3 above for its guaranteed IPTV service. When the first joiner of any session from any source say S9, joins a particular group, say G10, on a particular MAC domain, the CMTS selects the Downstream Channel Set to reach that joiner, creates GSF12 on that DCS and has GCR13 reference GSF12. The GCR 13 has a specific (S9, G10) criteria as GC 10 references GQC entry of type "Single-Session":

- GCR13: (S9,G10) IP DS (6,6,255) → GSF12.

When a joiner to a second session from any source, say S10, joins a particular group, say G11 on the MAC domain and the CMTS elects to distribute the session to the same DCS, the CMTS creates a separate GSF13 on that DCS and has GCR14 reference GSF13. The GCR 14 has a specific (S10, G11) criteria as GC10 references a GQC entry of type "Single-Session":

- GCR14: (S10,G11) IP DS (6,6,255) → GSF13.

IP Multicast packets forwarded by the CTMS to any MAC domain for session (S9,G10) will use GCR13 only if the IP DS field is set to 6 and hence will be transmitted using the appropriate GSF for that session. IP Multicast packets forwarded by the CMTS to any MAC domain for session (S9,G10) but with IP DS field not equal to 6 will instead be forwarded using the default GSF since no GCR is defined for those IP DS field values.

7.5.8.4 Default Group Service Flows

A CMTS MUST identify one of its Service Classes as the Default Group Service Class. When the CMTS replicates a multicast packet to a Downstream Channel Set on which the packet matches no Group Classifier Rule, the CMTS MUST transmit the packet on a Group Service Flow instantiated using the Default Group Service Class.

The CMTS MUST replicate unmatched IP multicast traffic only to a Downstream Channel Set that comprises an individual downstream channel. The CMTS does not replicate unmatched IP multicast traffic to downstream bonding groups. The Maximum Sustained Traffic Rate limit on the Default Group Service Class restricts the total amount of unclassified multicast traffic on each downstream channel. The CMTS MUST create a Default Group Service Flow on each of its downstream channels.

Because unmatched IP multicast traffic is required to be transmitted as non-bonded, the replication of a particular IP multicast session to a downstream bonding group requires the operator to either configure a GQC entry that matches the bonded multicast session or the CMTS to instantiate a GCR that matches the bonded multicast session in a vendor-specific manner.

7.5.8.5 Service Class QoS Parameter Changes

The CMTS MAY dynamically change the QoS parameters of all Group Service Flows derived from a Service Class when the QoS parameters of the Service Class are changed.

7.5.8.6 Group QoS Configuration Changes

Because the GC and GQC tables are the only mechanism for controlling the instantiation of GCRs and GSFs, when a GC or GQC entry is added, modified or deleted, the CMTS MUST dynamically implement changes to the GCR(s) and GSF(s) instantiated from that GC or GQC entry, as follows:

- For each replication of an IP multicast session on a DCS which matches a GC entry that references a valid GQC entry of type Aggregate, the CMTS MUST ensure that there exists a GCR that classifies the range of (S,G) from the matching GC entry to a GSF corresponding to the referenced aggregate type GQC entry.
- For each replication of an IP multicast session on a DCS which matches a GC entry that references a valid GQC entry of type Single, the CMTS MUST ensure that there exists a GCR that classifies the specific (S,G) onto a specific GSF corresponding to the referenced single type GQC entry.
- All GCRs which are not required to exist MUST be deleted.
- All GSFs which are not required to exist MUST be deleted.

NOTE 1: GCRs and GSFs may be created or deleted due to the following changes to the QoS configuration tables: adding a GC entry; deleting a GC entry; modifying the GC entry by changing the (S,G) range, the priority or other attributes; changing GQC entry reference; changing GQC QoS type, etc.

The time-frame for implementing changes to the GCRs and GSFs is not specified.

For Aggregated sessions the CMTS MUST assign sessions to a DCS such that number of sessions matching a GC entry referring to an Aggregate GQC entry does not exceed the Aggregate Session limit.

NOTE 2: Sessions may be dropped from a DCS by changing the AggregateSessionLimit and also perhaps due to the changes as noted above. The sessions which the CMTS chooses to keep or drop when the Aggregate Session limit is decreased are vendor-specific.

7.5.9 Other Multicast and Broadcast Traffic

The Group QoS Configuration Table specifies how QoS is provided to downstream multicast traffic only for joined IP Multicast sessions. The QoS provided to all other downstream broadcast and multicast traffic is not configured with the GQC Table.

Examples of traffic which are not configured with a GQC table include:

- locally generated IP multicasts (e.g. IP router advertisements such as RIPv2 and OSPF);
- DSG tunnel traffic;
- layer 3 broadcasts (e.g. DHCP broadcasts);
- layer 2 broadcasts (e.g. ARP); and
- layer 2 multicasts (e.g. Spanning Tree Protocol).

The CMTS MUST transmit and account for all layer 2 multicast and broadcast traffic with some Group Service Flow. The CMTS MAY define Group Classifier Rules that classify multicast and broadcast traffic other than for joined IP multicast traffic.

7.6 Downstream Traffic Priority

The downstream Traffic Priority parameter is an explicit tag that will allow the CM to support multiple prioritized egress queues at its CMCI port. DOCSIS 3.0 defines a Downstream Service Extended Header (DS EHDR) element (refer to clause 6.2.5.6) in which the first three bits of the EH_VALUE indicate the Traffic Priority of the packet. If the Traffic Priority takes the default value of 0, the CMTS is not required to include the DS EHDR on packets that do not require a DSID label.

The CM MUST support a minimum of two egress queues per CMCI port. The egress queue for a particular packet is selected by the Traffic Priority sub-element in the DS EHDR of the packet. If the DS EHDR is missing then the CM MUST assume the packet has the Default Priority of zero.

The CM MUST NOT transmit downstream packets of lower Traffic Priority while there are packets of higher Traffic Priority ready to transmit on the CMCI.

7.6.1 Traffic Priority Ordering and Mapping to CM Output Queues

Table 7-10 indicates the CM output queue to which a packet MUST be assigned based on the number of CM output queues supported by the CM implementation and the Traffic Priority indicated in the DS EHDR of the packet. The CM output queues are numbered in order of increasing priority with 0 as the lowest priority and 7 as the highest priority. If the DS EHDR is not present in the packet, a Traffic Priority of 0 is used.

Table 7-10: Mapping of Traffic Priority to CM output queue

		Number of CM output queues						
		2	3	4	5	6	7	8
Traffic Priority	0 (Default)	0	0	0	0	0	0	0
	1	0	0	0	0	0	0	1
	2	0	0	1	1	1	1	2
	3	0	0	1	1	2	2	3
	4	1	1	2	2	3	3	4
	5	1	1	2	3	4	4	5
	6	1	2	3	4	5	5	6
	7	1	2	3	4	5	6	7

7.7 Payload Header Suppression

The overview clause which follows explains the principles of Payload Header Suppression. The subsequent clauses explain the signalling for initialization, operation and termination. Finally, specific upstream and downstream examples are given. The following definitions are used.

Table 7-11: Payload Header Suppression Definitions

PHS	Payload Header Suppression	The suppressing of an original byte string at the sender and restoration of the suppressed byte string at the receiver.
PHSI-indexed PHS	PHSI-indexed Payload Header Suppression	Payload Header Suppression of unicast traffic, in either the upstream or downstream direction. The PHSI is used to identify the applicable PHS Rule.
DSID-indexed PHS	DSID-indexed Payload Header Suppression	Payload Header Suppression of multicast traffic in which the CMTS is the sender and the CM is the receiver. The DSID is used to identify the applicable PHS Rule.
PHSR	Payload Header Suppression Rule	A set of TLVs that apply to a specific identifier (PHSI or DSID) described in clause C.2.2.10.
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 uniquely references the PHSR for PHSI-indexed PHS.
PHSM	Payload Header Suppression Mask	A bit mask used to interpret the values in the PHSF; it indicates which bytes in the PHSF to suppress and which bytes to not suppress.
PHSS	Payload Header Suppression Size	The length of the Payload Header Suppression 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.
DSID	Downstream Service Identifier	A 20-bit value in the DS-EHDR which uniquely identifies the PHSR for DSID-indexed PHS.

7.7.1 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. 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. When a classifier is deleted, any associated PHS rule **MUST** also be deleted by the CM or CMTS.

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.

It is the responsibility of the higher-layer service entity to generate a PHSR 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.

There are two types of Payload Header Suppression, PHSI-indexed PHS and DSID-indexed PHS. PHSI-indexed PHS is Payload Header Suppression of unicast traffic in either the upstream or downstream direction and corresponds to DOCSIS 1.1/2.0 Payload Header Suppression. DSID-indexed PHS is Payload Header suppression of multicast traffic in the downstream direction.

7.7.1.1 PHSI-indexed PHS

In PHSI-indexed PHS, the MAC Extended Header contains a Payload Header Suppression Index (PHSI) which references the Payload Header Suppression Field (PHSF).

Although PHSI-indexed 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. PHSI-indexed 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.

The sending entity uses Classifiers to map packets into an Individual Service Flow. The Classifier uniquely maps packets to its associated Payload Header Suppression Rule. If the receiving entity is the CMTS, it uses the Service Flow ID and the PHSI to restore the PHSR. If the receiving entity is the CM, then the PHSI is sufficient 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 needs to be removed from the Service Flow and a new rule needs to be installed.

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 CM or CMTS MAY specify the PHSF and PHSS. This provision allows for pre-configured headers or for higher level signalling protocols outside the scope of the present document to establish cache entries.

PHS signalling of PHSI-indexed PHS via Registration or Dynamic Service messages allows for up to 255 Payload Header Suppression Rules per Individual Service Flow, but the exact number of PHS rules supported per Individual Service Flow is implementation dependent. Similarly, PHS signalling allows for PHSI-indexed PHS Sizes of up to 255 bytes; however, the maximum PHSI-indexed PHS Size supported is implementation dependent. For interoperability, the minimum PHS Size that MUST be supported by the CM and CMTS is 64 bytes for any PHSI-indexed 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.

7.7.1.2 DSID-indexed PHS

DOCSIS 3.0 introduces the feature of suppressing the payload headers of downstream DSID-labeled multicast traffic. If only the 8-bit PHSI was used to identify a PHS Rule on the CM, it would significantly limit the total number of multicast sessions on a MAC Domain that could perform Payload Header Suppression. For this reason, a mechanism is defined to utilize the 20-bit DSID identify a PHS Rule on a CM.

In DSID-indexed Payload Header Suppression, a repetitive portion of the payload headers following the MAC Destination Address and MAC Source Address is suppressed by the CMTS and restored by the CM. When using DSID-indexed PHS, the CMTS MUST ensure that the Payload Header Suppression Header contains an EH_VALUE of 255.

In DSID-indexed Payload Header Suppression, the CMTS is the sending entity. The CMTS uses Classifiers to map packets into a Group Service Flow which is not communicated to the CM. The CMTS uses a DSID to identify the associated Payload Header Suppression Rule. The CMTS MUST include a Downstream Service Extended Header containing the DSID with which the Payload Header Suppression Rule is associated when performing DSID-indexed PHS. The CM uses the DSID to restore the PHSR. The CM MUST use the DSID to identify the PHS Rule to be used when unsuppressing the packet when the Payload Header Suppression Header contains an EH_VALUE of 255.

In order to modify a DSID-indexed PHS Rule, the CMTS removes the old rule from the DSID and installs a new rule using the DBC messages. The CMTS is responsible for ensuring that all CMs receiving the multicast stream labeled with the DSID are aware of the modified DSID-indexed PHS Rule. Whenever there is a change to the PHS Rule configured for a session then the CMTS SHOULD signal the required PHS rules using DBC messages to all the CMs which are forwarding that Multicast session. When a DSID is deleted, both the CM and the CMTS MUST delete any associated PHS rule.

The CMTS MUST NOT implement DSID-indexed PHS on a Multicast session when there are CMs forwarding that session which are not capable of DSID-indexed PHS.

PHS signalling via Registration or Dynamic Bonding Change messages allows for one DSID-indexed Payload Header Suppression Rule per multicast DSID. The CM MUST support at least 16 DSID-indexed PHS rules. The minimum PHS Size that MUST be supported by a CM or CMTS is 64 bytes for any DSID-indexed PHS rule. The CMTS SHOULD NOT signal the CM a DSID-indexed PHS rule with a PHS size greater than 64 bytes. As with any other parameter requested in a Dynamic Bonding Change, a PHS-related DBC request can be denied because of a lack of resources.

7.7.2 Example Applications

Assume 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 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 PHSI lookup at the receiving entity will always yield the correct result.

Assume 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.

A DSID uniquely identifies a multicast video stream. The CMTS classifies the multicast video stream to a group service flow, communicating the DSID, the DSID's multicast attributes and the associated PHS rule to the CM. The DSID has a PHSR, which suppresses the Protocol Type of UDP, static fields in the IPv6 header and static fields in the UDP header. The CMTS suppresses the header of this video stream according to the PHSR. The CMTS then adds a 2-byte Payload Suppression Header containing the EH_VALUE of 255 to every packet in that video stream.

7.7.3 Operation

To clarify operational packet flow, this clause 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 clause is followed. Figure 7-18 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 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 Payload Header Suppression EHDR Sub-Element, clause 6.2.5.4 and queue the packet on the Upstream Service Flow.

When the packet is received by the CMTS, the CMTS will determine the associated Service Flow either by internal means or from other Extended Headers elements. The CMTS uses the Service Flow 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 bytes. If verification was not enabled, then there is no guarantee that the PHSF bytes will match the original header bytes.

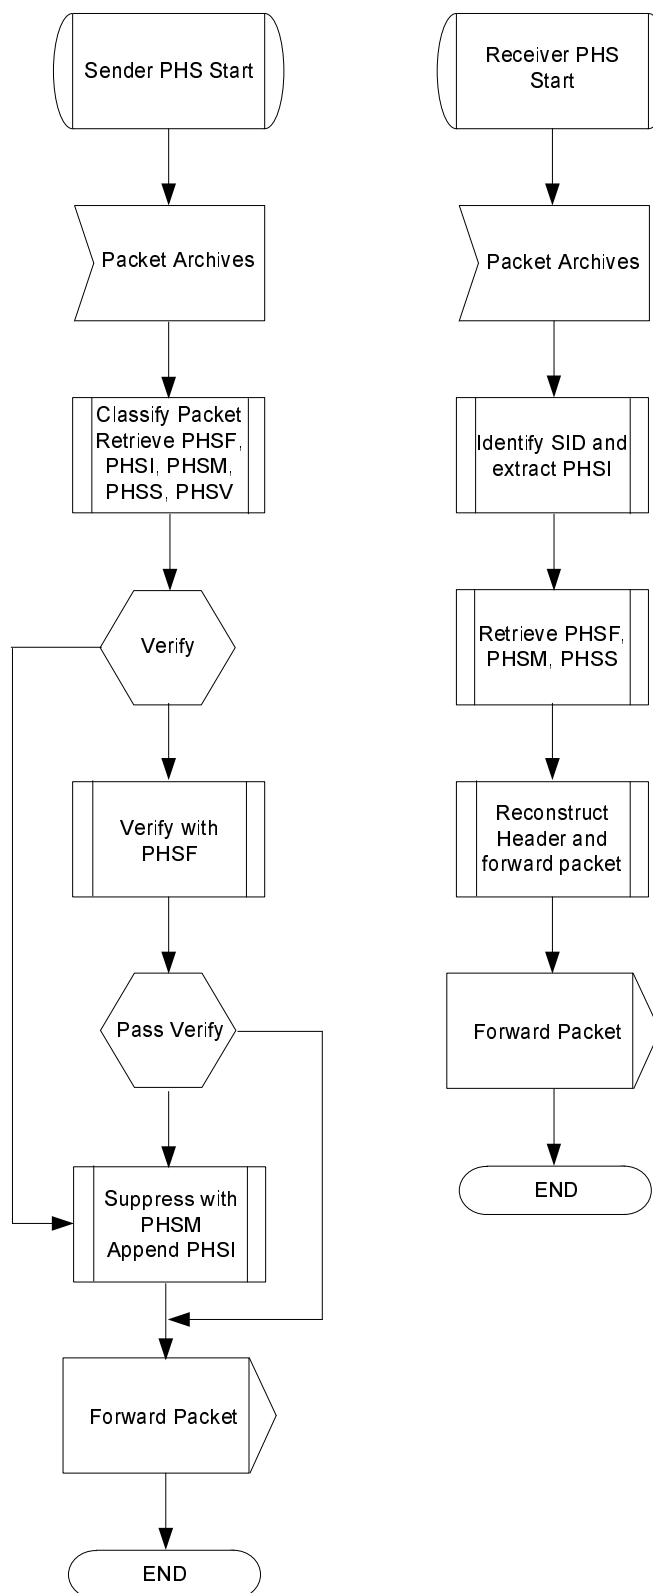


Figure 7-18: Payload Header Suppression Operation

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 Payload Header Suppression EHDR Sub-Element field of the Service Flow EH Element and queue the packet on the Downstream Service Flow. The Ethernet header can not be suppressed on the downstream because it is needed for MAC filtering on the CM.

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 expands the packet and then proceeds with normal packet processing.

NOTE 1: On the upstream there are up to 255 PHSI per service flow (and multiple service flows per CM); while on the downstream there are 255 PHSI per CM. The higher granularity on the upstream is useful for the multiple grants per interval feature.

Figure 7-19 demonstrates packet suppression and restoration when using PHS masking. Masking allows only bytes which do not change to be suppressed.

NOTE 2: The PHSF and PHSM span the entire Suppression Field, including suppressed and unsuppressed bytes.

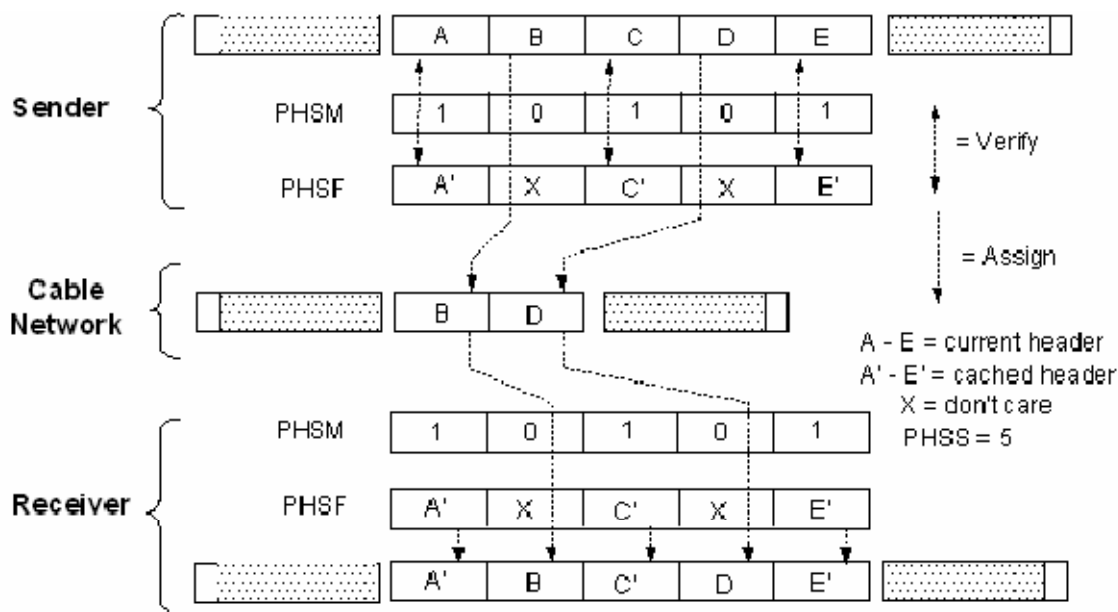


Figure 7-19: Payload Header Suppression with Masking

7.7.4 Signalling

7.7.4.1 Signalling PHSI-Indexed Payload Header Suppression

PHSI-indexed Payload Header Suppression requires the creation of three objects:

- Individual Service Flow.
- Classifier.
- Payload Header Suppression Rule.

These three objects MAY be created in separate message flows by the CM or CMTS or be created simultaneously.

PHSI-indexed PHS Rules are created with Registration, DSA or DSC messages. The CMTS MUST define the PHSI when the PHSI-indexed 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 7-20 shows the two ways CMTS initiated or CM initiated, to signal the creation of a PHSI-indexed PHS Rule on an Individual Service Flow.

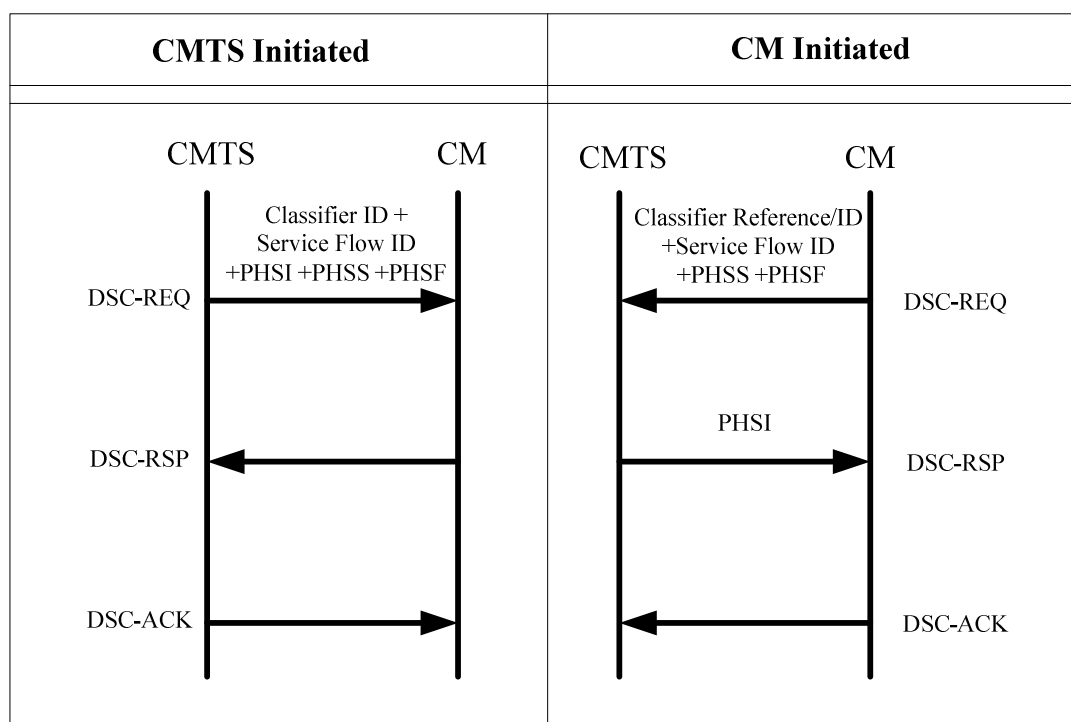


Figure 7-20: PHSI-Indexed Payload Header Suppression Signalling Example

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

As an example, it is likely that when an Individual 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 PHSI-indexed 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 signalling by the CM or CMTS. PHSV and PHSM fields have default values (see clauses C.2.2.10 and C.2.2.10.3) and 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 by the CM or CMTS via DSC signalling.

Each step of the PHS rule definition performed by the CM or CMTS, 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 Flow 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.

7.7.4.2 Signalling DSID-Indexed Payload Header Suppression

DSID-Indexed Payload Header Suppression of multicast traffic requires that the CMTS signal two objects to the CM:

- DSID with multicast attributes.
- Payload Header Suppression Rule.

These two objects **MAY** be signaled in separate message flows by the CMTS or be signaled simultaneously.

The CMTS defines the DSID, the PHSS and the PHSF when it creates the DSID-indexed PHS Rule and signals these parameters to the CM via Registration or DBC messages. DSID-Indexed PHS Rules are deleted with the DBC message.

Figure 7-21 demonstrates signalling of the creation of a DSID-Indexed PHS Rule.




Figure 7-21: Signalling of a DSID-Indexed Payload Header Suppression Rule

Unlike PHSI-indexed PHS, a DSID-Indexed PHS rule **MUST** be fully defined by the CMTS at the time a DSID-indexed PHS rule is signaled to the CM. A fully defined DSID-indexed PHS rule includes a DSID with multicast attributes and both the PHSF and PHSS field values. If PHSV and PHSM are not known when the DSID-indexed PHS rule is signaled, the CM uses their default values (see clauses C.2.2.10 and C.2.2.10.3). There is no mechanism for the CMTS to modify DSID-indexed PHS rules. The CMTS may add new DSID-indexed PHS rules or delete existing DSID-indexed PHS rules, but may not modify DSID-indexed PHS rules.


7.7.5 Payload Header Suppression Examples

7.7.5.1 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 [i.8] VoIP traffic in the upstream with Unsolicited Grant Service. When Classifiers are added to the flow, the classifier contains a PHSR with a PHSS value of 42, a PHSM and a PHSV which explicitly states that the first 42 bytes following the MAC Extended Header on all packets in that flow are 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 clause C.2.2.6.6).

Figure 7-22 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.

a) VoIP with Normal Encapsulation



b) VoIP with Header Suppression

Figure 7-22: Upstream Payload Header Suppression Example

Figure 7-22a 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 mini-slots. 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.

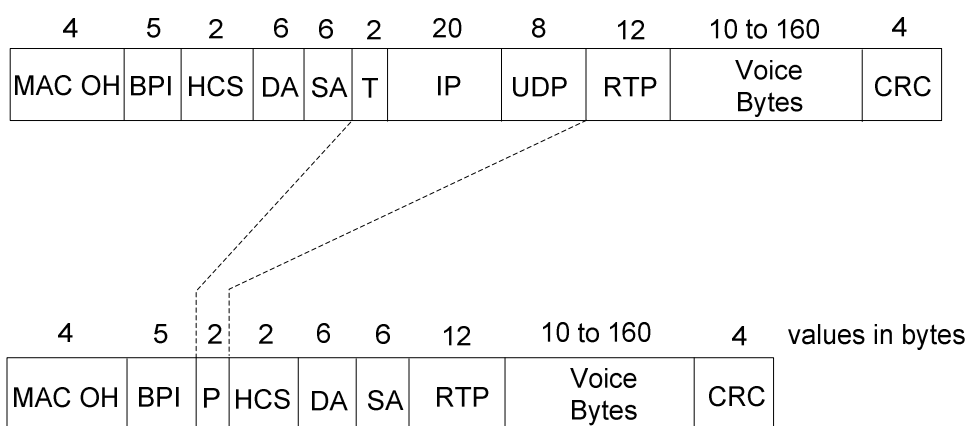
Figure 7-22b shows the same payload with Payload Header Suppression enabled. In the upstream, Payload Header Suppression begins with the first byte after the MAC Header Checksum. The 14 byte Ethernet header, the 20 byte IP header and the 8 byte UDP header have been suppressed and a 2 byte PHS Extended 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.

7.7.5.2 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 are 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 7-23 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.

a) VoIP with Normal Encapsulation



b) VoIP with Header Suppression

Figure 7-23: Downstream Payload Header Suppression Example


Figure 7-23a 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 and one of the DOCSIS 3.0 header types (not included for pre-DOCSIS 3.0), 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 7-23b 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.

7.7.5.3 DSID-Indexed Multicast Example

A multicast video stream is established. When the multicast DSID is created to label the multicast video stream, a PHSR with a PHSS value of 48 and a PHSM which explicitly indicates that 45 bytes of the payload header on all packets are processed for suppression and restoration is included. Any packet which does not pass PHSV correctly will not have its header suppressed but will be transmitted subject to the traffic shaping rules in place for that DSID.

a) VOD with Normal Encapsulation



b) VOD with Header Suppression

Figure 7-24: DSID-Indexed Payload Header Suppression Example

Figure 7-24 shows the encapsulation used in the downstream with and without DSID-Indexed Payload Header Suppression. An RTP Payload without IPsec is used as a specific example.

Figure 7-24a shows a normal multicast video packet transmitted to a CM. The Layer 2 overhead includes the 6 byte MAC header with a 6 byte DS Extended Header (which contains the 20-bit DSID used to identify the PHS rule), 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 multicast video payload uses a 40-byte IPv6 header, an 8 byte UDP header and a 12-byte RTP header.

Figure 7-24b 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, 35 bytes of the 40-byte IP header and 6 bytes of the 8-byte UDP header have been suppressed. A 2 byte PHS Extended Header element with an EH_VALUE of 255 has been added, for a net reduction of 41-bytes. In this example of an established multicast video stream, these fields remain constant from packet to packet and are thus redundant.

7.8 Data Link Encryption Support

The procedures to support data link encryption are defined in [15]. The interaction between the MAC layer and the security system is limited to the items defined below.

7.8.1 MAC Messages

MAC Management Messages (clause 7.8.1) **MUST NOT** be encrypted, except for certain cases where such a frame is included in a Pre-3.0 DOCSIS fragmented concatenated burst on the upstream (refer to clause 7.8.3). For Multiple Transmit Channel Mode operation, MAC Management Messages **MUST NOT** be encrypted.

7.8.2 Framing

When encryption is applied to a data PDU, the CM **MUST** include the Privacy EH element [15] as the first EH element of the Extended Header field (EHDR). When encryption is applied to a data PDU, the CMTS **MUST** include the Privacy EH element [15] as the first EH element of the Extended Header field (EHDR).

7.8.3 Multiple Transmit Channel Mode Operation and Packet Encryption

For Multiple Transmit Channel Mode Operation, when enabled for a service flow, encryption **MUST** be performed on data PDUs prior to Continuous Concatenation and Fragmentation at the CM. At the CMTS, packets **MUST** be reassembled prior to any decryption.

8 Channel Bonding

Channel bonding refers to the scheduling of information in DOCSIS service flows over multiple channels concurrently. In the downstream direction, the CMTS distributes individual packets over multiple channels and usually includes a Downstream Service Extended Header that contains a packet sequence number that permits the CM to resequence out-of-order packets. In the upstream direction, the CM continuously concatenates and fragments a stream of packets into a set of "segments" and distributes those segments over the grants scheduled by the CMTS for the service flow. Each segment has a sequence number to permit the CMTS to re-order segments received out of order. A service flow which has information scheduled over multiple channels is called a "bonded" service flow. A set of two or more channels over which the CMTS schedules the information of a service flow is called a "bonding group".

8.1 Upstream and Downstream Common Aspects

8.1.1 Service Flow Assignment

The CMTS MUST assign Service Flows to either individual upstream or downstream channels or to upstream or downstream bonding groups. This assignment can be dynamic in that, at any point in time, the CMTS can reassign a Service Flow to a different channel or bonding group following the guidelines in this clause.

When a service flow is assigned to a bonding group, the CMTS MUST assign the service flow to all channels of the bonding group. When a service flow with resequencing enabled is assigned to a downstream bonding group, the CMTS MUST label the packets of the service flow with a DSID whose Resequencing Channel List is set to contain all channels of the bonding group. When a service flow is assigned to an upstream bonding group, the CMTS MUST assign SIDs for all channels of the bonding group. These requirements apply to the administrative assignment of service flows to bonding groups and are not intended to imply requirements on the CMTS scheduler, e.g. the CMTS is not required to schedule traffic on all channels of the bonding group.

DOCSIS 3.0 introduces the concept of assigning Service Flows to channels or bonding groups based on binary attributes. Some of these binary attributes are defined below, while others are left for operator definition. The specification-defined attributes have specific default values based on the characteristics of the channel or bonding group. The operator-defined attributes default to zero. The operator can configure a Provisioned Attribute Mask for each channel and provisioned bonding group to assign values for the operator-defined binary attributes and/or to override the default values of the specification-defined attributes. The operator may configure, in the CM configuration file, a Required Attribute Mask and a Forbidden Attribute Mask for a service flow. Additionally, in a CM-Initiated Dynamic Service Request, the CM could include a Required Attribute Mask and a Forbidden Attribute Mask for a service flow. The CMTS attempts to assign service flows to channels or bonding groups such that all required attributes are present and no forbidden attributes are present. The attribute based assignment applies both to the initial assignment of the Service Flow, as well as to any subsequent reassignment. Attribute-based assignment applies to both Individual Service Flows and Group Service Flows. The CMTS may use other mechanisms for assigning service flows to channels or bonding groups, such as the application ID or other service flow parameters.

The Cable Operator determines a set of attributes of interest that can be applied to an upstream or downstream channel or bonding group. Examples of binary attributes of a downstream interface include:

- bonded, whether or not the downstream interface represents a bonding group;
- high availability, e.g. the existence of spare hardware that can automatically take over for a failed channel;
- M-CMTS, whether the channel is an M-CMTS DEPI tunnel or an integrated RF channel;
- low latency, e.g. whether the channel has a lower than usual latency due to a lower interleaver delay;
- DSG, i.e. intended as a single downstream channel on which to put all DSG CMs;
- IPVideo, i.e. intended as a DBG on which to put all IP Video;
- business, i.e. intended for business committed information rate service; and
- synchronized, i.e. whether the channel is synchronized to the upstream master clock.

Examples of upstream interface attributes are:

- bonded, whether or not the upstream interface represents a bonding group;
- high availability; e.g. the existence of spare hardware that can automatically take over for a failed channel;
- low latency, e.g. whether the channel has a lower than usual latency due to CMTS scheduling policy;
- high robustness, e.g. modulation and/or FEC parameters that provide for low packet error rate.

Associated with each channel or provisioned bonding group is a "Provisioned Attribute Mask" with a 1 or 0 in each bit position of a 32-bit integer. The Attribute Masks follow the BITS Encoding convention where the most significant bit of the Mask is considered bit 0. The specification-defined attributes are bits 0 through 15 of the Attribute masks. The remaining bits are left for operator-definition.

To assist with initial deployments of DOCSIS 3.0, to simplify configuration and in order to allow for consistent configurations across different vendor CMTSs, the specification-defined attribute bits and their default values are defined below.

Bit position (0): bonded.

Resource	Default value
DS channel	The CMTS MUST set this bit to zero for all individual Downstream Channels.
DSBG	The CMTS MUST set this bit to one for all Downstream Bonding Groups.
US channel	The CMTS MUST set this bit to zero for all individual Upstream Channels.
USBG	The CMTS MUST set this bit to one for all Upstream Bonding Groups.

Bit position (1): Low Latency.

Resource	Default value
DS channel	The CMTS SHOULD set this bit to one when the corresponding channel is configured to provide relatively low latency service.
DSBG	The CMTS SHOULD set this bit to one when all channels in the bonding group provide relatively low latency service and the CMTS can communicate a DSID Resequencing Wait Time less than the maximum DSID Resequencing Wait Time (see annex B).
US channel	The CMTS SHOULD set this bit when a channel provides relatively low latency service.
USBG	The CMTS SHOULD set this bit to one when all channels in the bonding group provide relatively low latency service.

The term "relatively low latency service" is left for vendor definition.

Bit position (2): High Availability.

Resource	Default value
DS channel	The CMTS SHOULD set this bit to one when the corresponding channel provides High Availability features.
DSBG	The CMTS SHOULD set this bit to one when all of the corresponding channels provide High Availability features.
US channel	The CMTS SHOULD set this bit to one when the corresponding channel provides High Availability features.
USBG	The CMTS SHOULD set this bit to one when all of the corresponding channels provide High Availability features.

The definition of what constitutes "High Availability features" is vendor-specific.

Bit positions (3..15): reserved for future use (default value 0).

Each Service Flow is optionally configured with the following TLV parameters:

- Service Flow Required Attribute Mask;
- Service Flow Forbidden Attribute Mask; and
- Service Flow Attribute Aggregation Rule Mask.

When present in a Service Flow encoding in the CM configuration file, these TLVs are sent in the Registration Request. These parameters also could be present in a Dynamic Service message originated by the CM. When these parameters are not present in the service flow encoding, then attribute-based assignment does not apply and the CMTS may assign the flow to a channel or bonding group as it sees fit.

Attribute based assignment means that the CMTS assigns Service Flows to interfaces such that all required attributes are present and all forbidden attributes are absent.

In the case of assignment to an individual upstream or downstream channel, the CMTS assigns a Service Flow to a channel for which all of the Required attributes are present and all the Forbidden attributes are absent. The Service Flow Attribute Aggregation Rule Mask is ignored. When assigning a Service Flow to an individual channel and a Required Attribute Mask is defined for a Service Flow, the CMTS MUST assign the Service Flow to a channel which has a 1 bit in all positions of its Provisioned Attribute Mask corresponding to 1 bits in the Service Flow Required Attribute Mask, if such a channel is available to be included in the CM's Receive Channel Set. When assigning a Service Flow to an individual channel and a Forbidden Attribute Mask is defined for a Service Flow, the CMTS MUST assign the Service Flow to a channel which has a 0 bit in all positions of its Provisioned Attribute Mask corresponding to a 1 bit in the Forbidden Attribute Mask, if such a channel is available to be included in the CM's Receive Channel Set. If no channel is available which satisfies the Service Flow Required Attribute Mask and Service Flow Forbidden Attribute Mask for the Service Flow, the CMTS is free to assign the Service Flow to any channel in the MD-CM-SG of the CM.

In the case of assignment to a provisioned upstream or downstream bonding group, the operation is identical to the case of assignment to an individual upstream or downstream channel. The CMTS assigns a Service Flow to a bonding group for which all of the Required attributes are present and all the Forbidden attributes are absent. The Service Flow Attribute Aggregation Rule Mask is ignored. When assigning a Service Flow to a provisioned bonding group and a Required Attribute Mask is defined for a Service Flow, the CMTS MUST assign the Service Flow to a bonding group which has a 1 bit in all positions of its Provisioned Attribute Mask corresponding to 1 bit in the Service Flow Required Attribute Mask, if such a bonding group is available to be included in the CM's Receive Channel Set. When assigning a Service Flow to a provisioned bonding group and a Forbidden Attribute Mask is defined for a Service Flow, the CMTS MUST assign the Service Flow to a bonding group which has a 0 bit in all positions of its Provisioned Attribute Mask corresponding to a 1 bit in the Forbidden Attribute Mask, if such a bonding group is available to be included in the CM's Receive Channel Set. If no bonding group is available which satisfies the Service Flow Required Attribute Mask and Service Flow Forbidden Attribute Mask for the Service Flow, the CMTS is free to assign the Service Flow to any bonding group in the MD-CM-SG of the CM. Alternatively, the CMTS could dynamically create a bonding group which satisfies the Attribute Masks for the Service Flow.

The CMTS MAY support the dynamic creation of upstream or downstream bonding groups. In the case of assignment to a dynamically created upstream or downstream bonding group, the CMTS MUST assign a Service Flow to a Dynamic Bonding Group based on the values of the Service Flow Attribute Aggregation Rule Mask and the Provisioned Attribute Masks of the individual channels of the bonding group. To perform the comparison, the bits corresponding to a particular attribute on all candidate channels are logically combined via either an AND operation or an OR operation, depending on the setting of the Service Flow Attribute Aggregation Rule Mask. The exception to this is the "bonded" attribute bit, for which the result of the combination is defined to always be 1 (regardless of the setting of the Service Flow Attribute Aggregation Rule Mask). The result of the combination is then compared with the Service Flow Required Attribute Mask and the Service Flow Forbidden Attribute Mask. If the Service Flow Required Attribute Mask has a 1 in a particular bit position, the CMTS MUST assign the Service Flow to a Bonding Group for which the combination result also has a 1 in the corresponding bit position. If the Service Flow Forbidden Attribute Mask has a 1 in a particular bit position, the CMTS MUST assign the Service Flow to a Bonding Group for which the combination result has a 0 in the corresponding bit position. If no dynamic bonding group can be created and no existing bonding group is available to satisfy the Service Flow Required Attribute Mask and Service Flow Forbidden Attribute Mask for the Service Flow, the CMTS is free to dynamically create any bonding group or assign the Service Flow to any existing bonding group (provisioned or dynamically created) in the MD-CM-SG of the CM.

If the CMTS does not assign a Service Flow such that Required and Forbidden Attributes are met, it MUST log an event and update the MIB to report the attribute assignment failure. If a CMTS configuration change results in Service Flows being assigned to channels or bonding groups that do not match their Required and Forbidden Attributes, the CMTS MUST log an event and update the MIB to report the mismatch.

The operator is responsible for defining the Provisioned Attribute Mask for provisioned bonding groups. In particular, the operator is responsible for interpreting how the attributes of individual interfaces aggregate to the attribute of the bonding group. For example, a bonding group may be configured with a "High Availability" attribute only when all of its component channels have "High Availability", but a bonding group may also be configured with "High Latency" when any of its channels have "High Latency".

Although the attributes are defined as binary values, an attribute mask bit position may represent a particular range of a variable. For example, one attribute bit position may represent the attribute "Intended for maximum rates exceeding 50 Mbps" and only bonding groups with sufficient capacity to meet that maximum rate will have that attribute set in the bonding group's Provisioned Attribute Mask.

The following table summarizes the CMTS assignment of rules for the various combinations of corresponding bits in the Service Flow Required Attribute Mask, the Service Flow Forbidden Attribute Mask and the Service Flow Attribute Aggregation Rule Mask for dynamically created bonding groups.

Table 8-1: Attribute Mask Summary Table for Attribute Bits Other than the "Bonded" Attribute

SF Required Attribute Mask	SF Forbidden Attribute Mask	SF Attribute Aggregation Rule Mask (1=AND, 0=OR)	Interpretation
0	0	0	Do not care
0	0	1	Do not care
0	1	0	No channels can have this attribute turned on (default if Forbidden bit is set and Rule is unspecified)
0	1	1	At least one channel has this attribute turned off
1	0	0	At least one channel has this attribute turned on
1	0	1	All channels have this attribute turned on (default if Required bit is set and Rule is unspecified)
1	1	0	Not allowed
1	1	1	Not allowed

Table 8-2: Attribute Mask Summary Table for the "Bonded" Attribute Bit

SF Required Attribute Mask	SF Forbidden Attribute Mask	SF Attribute Aggregation Rule Mask (1=AND, 0=OR)	Interpretation
0	0	X	This Service Flow can be assigned to an individual channel or to a bonding group.
1	0	X	This Service Flow can be assigned to a bonding group (static or dynamic).
0	1	X	This Service Flow cannot be assigned to a bonding group (static or dynamic).
1	1	X	Not allowed.

NOTE: The Service Flow Attribute Aggregation Rule Mask does not apply to the "bonded" attribute bit. The Service Flow Required Attribute Mask and Service Flow Forbidden Attribute Mask directly control whether the service flow is assigned to a bonding group (static or dynamic) or to an individual channel.

8.1.2 CMTS Bonding and Topology Requirements

The CMTS MUST permit Downstream Channels reaching the same CM-SG to be configured into separate MAC Domains. The CMTS MUST permit Upstream Channels reaching the same CM-SG to be configured into separate MAC Domains. This permits an operator to segregate tiers of service (e.g. DSG CMs or business service CMs) to entirely separate MAC Domains.

The CMTS MUST enforce that Downstream RF Channels reaching the same CM-SG are configured to different frequencies.

The CMTS MUST enforce that Upstream physical Channels reaching the same CM-SG are assigned to different frequencies.

The CMTS MUST enforce that all Downstream Channels in a Downstream Bonding Group are from the same MAC Domain. The CMTS MUST enforce that all Upstream Channels in an Upstream Bonding Group are from the same MAC Domain.

The CMTS MUST support provisioned downstream bonding groups containing 2 to 4 channels. The CMTS MAY support provisioned downstream bonding groups containing more than 4 channels. The CMTS MAY support dynamically created downstream bonding groups.

The CMTS MUST support provisioned upstream bonding groups containing 2 to 4 channels. The CMTS MAY support provisioned upstream bonding groups containing more than 4 channels. The CMTS MAY support dynamically created upstream bonding groups.

To efficiently utilize downstream bandwidth across cable modems with different receive channel capabilities and/or bonding groups of different sizes, the operator may wish to assign individual downstream channels to multiple, overlapping Downstream Bonding Groups. With this configuration, when a channel is associated with multiple bonding groups, its bandwidth is available for use by the CMTS to carry traffic for any of the bonding groups with which it is associated.

While the CMTS is expected to manage bandwidth efficiently over overlapping bonding groups, it should be recognized that managing bandwidth in this configuration is unique to cable and may require the use of complex algorithms, especially when the number of overlapping bonding groups becomes large. For this reason, this specification places no requirements on how the CMTS should allocate the channel bandwidth among multiple overlapping bonding groups. A CMTS vendor may choose an algorithm that simplifies the scheduling, load balancing and management of overlapping bonding channels by placing vendor-specific limitations on the bonding group to channel assignment.

The CMTS **SHOULD** support the ability to include each downstream channel in at least four provisioned downstream bonding groups simultaneously. The CMTS **MAY** support the ability to include the same downstream channel in more than four provisioned downstream bonding groups simultaneously.

8.2 Downstream Channel Bonding

8.2.1 Multiple Downstream Channel Overview

With DOCSIS 2.0, downstream data service is provided to a Cable Modem on a single downstream channel. DOCSIS 3.0 significantly expands the downstream service offering by requiring DOCSIS 3.0 CMs to be capable of receiving multiple downstream channels simultaneously.

DOCSIS 3.0 allows the CMTS to assign individual downstream service flows to particular downstream channels. For example, a CMTS may assign a video-over-IP service flow to a downstream channel with deeper interleaving for higher reliability, while also assigning a VOIP flow destined for the same modem to a different downstream channel with shallower interleaving for low latency.

DOCSIS 3.0 also supports the concept of "Downstream Channel Bonding", in which independent streams of packets are distributed across the multiple downstream channels of a Downstream Bonding Group. Downstream Channel Bonding allows a DOCSIS 3.0 CM to forward data at greater than the throughput of a single downstream channel. Downstream Channel Bonding can reduce the delay of individual downstream packets. Downstream Channel Bonding can reduce the admission failures of large-bandwidth flows like HDTV by allowing the flow to share bandwidth across multiple downstream channels, rather than having to be admitted completely to a single channel.

The CMTS makes the decision whether to assign each downstream service flow either to a bonding group or to a single downstream channel. A downstream service flow assigned to a bonding group is called a "downstream bonded service flow". A downstream service flow assigned to a single channel is called a "downstream non-bonded service flow". The CMTS is free to assign some downstream service flows as bonded and some service flows as non-bonded. The CMTS is free to change the scheduling of a given downstream service flow between bonded and non-bonded, although certain requirements apply for communicating the channel set for sequenced packets to the CM.

With bonded service flows, the CMTS transmits the packets onto the multiple channels of a Downstream Bonding Group. The CMTS transmits each complete packet on a single channel. By default, packets of a bonded service flow are sequenced in order to guarantee in-order forwarding by the CM. In the absence of explicit, vendor-specific configuration to the contrary, the CMTS **MUST** transmit the packets of each bonded Service Flow with a 5-byte DS EHDR. The CMTS **MAY** support a vendor-defined configuration option to schedule certain service flows, e.g. for VOIP, as distributed over the multiple channels of a bonding group without sequencing the packets. When this option is applied, the order in which packets received on different downstream channels are forwarded by the CM is not guaranteed.

The CMTS **MAY** sequence the packets of non-bonded service flows; this can prevent out-of-order delivery when moving a service flow to a different channel for load balancing purposes.

8.2.2 CMTS Downstream Bonding Operation

A Downstream Bonding Group is a set of Downstream Channels on which the CMTS distributes packets. Downstream Bonding Groups may either be statically configured or dynamically determined by the CMTS. The CMTS **MUST** support the static configuration and modification of Downstream Bonding Groups. The CMTS **MAY** support the dynamic creation and/or modification of Downstream Bonding Groups.

To facilitate resequencing operations, the CMTS communicates to the CM a Downstream Resequencing Channel List for each Resequencing DSID. The Downstream Resequencing Channel List contains a list of channels on which the CM receives packets labeled with that DSID. In many cases it is identical to the channels in a Downstream Bonding Group. If there is no Downstream Resequencing Channel List for a Resequencing DSID, the CM receives packets labeled with that DSID on any channel in the Receive Channel Set. If the CMTS explicitly communicates a Downstream Resequencing Channel List for a Resequencing DSID to the CM, the CMTS **MUST** limit distribution of packets labeled with that DSID to the channels in the Downstream Resequencing Channel List. If the CMTS does not explicitly communicate a Downstream Resequencing Channel List for a Resequencing DSID the CMTS **MUST** distribute packets labeled with that DSID on the channels in the Receive Channel Set of CMs receiving that DSID.

The CMTS **MAY** dynamically change the assignment of a Service Flow to a different Downstream Channel or Bonding Group at any time. The CMTS **MAY** change a downstream Service Flow's assignment without notifying the CM(s) as long as the new channels are included in the Downstream Resequencing Channel List of the Resequencing DSID used for the packets of the Service Flow.

The CMTS **MUST** enforce that all Downstream Channels of a Downstream Bonding Group are contained within the same MAC Domain Downstream Service Group. A CMTS **MUST** permit configuration of a Downstream Channel as a member of multiple Downstream Bonding Groups. A CMTS **MAY** restrict the assignment of Downstream Channels to specific Downstream Bonding Groups based on vendor product implementation. For example, a CMTS product implementation may restrict the set of Downstream Channels that may be bonded in a given Downstream Bonding Group to only the subset of channels on a single line card.

8.2.3 Sequenced Downstream Packets

When packets are transmitted with a Resequencing DSID, they are called "sequenced" downstream packets. A CMTS transmits sequenced downstream packets with a five-byte Downstream Service Extended Header (DS EHDR). Each DS EHDR of a sequenced downstream packet defines the following fields relevant to the resequencing operation (clause 6.2.5.6, 5 byte EHDR):

- a 20-bit Downstream Service ID (DSID);
- a 1-bit Sequence Change Count; and
- a 16-bit Packet Sequence Number.

The DSID and the Sequence Change Count define a number space of Packet Sequence Numbers. The Packet Sequence Number identifies a packet's position within a sequence.

Ideally, the CMTS would always transmit packets in order of increasing Packet Sequence Number (i.e. it would always send a higher-numbered packet after or simultaneously with a lower-numbered packet, regardless of which channel(s) the packets are being released on). In practice, the CMTS cannot precisely meet this goal, so it is allowed to send higher-numbered packets earlier than lower-numbered packets on different channels by some amount (specified below). However, on any individual channel, the CMTS always transmits sequenced packets in order of increasing sequence number.

The CMTS **MUST** transmit sequenced downstream packets with a Resequencing DSID (see clause C.1.5.4.3.1) signaled to the CM or CMs intended to forward the sequenced packets.

A CMTS **MAY** initially use either Sequence Change Count zero (0) or one (1) in the DS EHDR of a newly created Resequencing DSID. The CMTS **MUST** use a Packet Sequence Number of zero (0) in the DS EHDR of the first packet transmitted on a newly created Resequencing DSID.

The CMTS MAY change the Sequence Change Count with any packet in a sequence. The CMTS MUST continue to transmit with the same Sequence Change Count for at least the Sequence Hold timeout (annex B). The CMTS MUST use a Packet Sequence Number of zero (0) in the DS EHDR of the first packet transmitted with the new Sequence Change Count. After receiving a sequenced packet with a new Sequence Change Count a CM MAY discard sequenced packets with the previous Sequence Change Count for a period no longer than the Sequence Hold timeout defined in annex B. If a packet is received after the expiration of the Sequence Hold timeout with the alternate Sequence Change Count, the CM MUST consider it to be another change event.

8.2.3.1 Downstream Sequencing

Once released from the CMTS, packets may experience varying delays before reaching the CM. This is particularly true in an M-CMTS system, where the packet traverses the CIN and EQAM. Packets are assumed to remain in order within a particular downstream channel, but packets on different channels may experience different delays. Hence, by the time packets arrive at the CM, lower-numbered packets may have been further delayed relative to higher-numbered packets on different channels. The amount of time that a higher-numbered packet is received earlier than a lower-numbered packet is called "skew" and is described in detail in clause 8.2.3.2.

The CM is responsible for receiving the packets of the stream on its multiple downstream channels, then putting packets back in the proper order as indicated by the Packet Sequence Numbers. This operation is termed "resequencing". Because packets may be received out of order across channels, the CM will have to be prepared to store higher-numbered packets for some amount of time while waiting for lower-numbered packets to arrive. The amount of storage needed depends on the expected skew at the CM, the expected packet rate and burstiness and other application-specific parameters. This specification sets limits on the amount of skew the CM should expect, but the other parameters are outside of the scope of the present document. Hence, the amount of storage provided by the CM for resequencing is vendor-specific.

Since the CM's storage space is limited, at any given moment it will have a limited range of sequence numbers it can consider "in range" as defined below. On occasion, the CM may receive one or more "out of range" sequence numbers. This could occur due to PHY-layer errors or bursts of errors, a temporary "excessive skew" event in the path between CMTS and CM or other reasons. The CM MUST discard these packets. The CM discards these packets in order to have enough room to store packets with in-range sequence numbers. If the CM has not received an "in range" packet for more than two minutes for a particular DSID and has discarded more than 1 000 "out of range" packets for that DSID, the CM MUST discard the current Next Expected Packet Sequence Number and attempt to establish a new value for the Next Expected Packet Sequence Number based on actual received Packet Sequence Numbers.

When the CM discards an "out of range" packet, it prepares a CM-STATUS message to indicate the event. If an "in range" packet is received prior to sending the CM-STATUS message, the CM does not transmit the message. This is described in clause 6.4.34.

A CM may be asked to perform resequencing on more than one stream of packets at a time. Each stream is identified by a DSID, clause 7.4. Packet Sequence Numbering is per-DSID and packets with different DSIDs may arrive and/or be forwarded by the CM in any order relative to each other. Thus, the CM operates a fully independent resequencing context and associated state machine for each DSID. As described in clause 7.4, the CM is required to support at least 16 resequencing contexts.

All mathematical operations on Packet Sequence Number are defined to be unsigned and modulo the field size (i.e. modulo 2^{16}). In particular, modulo arithmetic is used when comparing two Packet Sequence Numbers. A 16-bit value A is greater than a 16-bit value B if $[(A - B) \bmod 2^{16}] < 2^{15}$. A 16-bit value A is less than a 16-bit value B if $[(A - B) \bmod 2^{16}] \geq 2^{15}$.

Packet Sequence Numbers and Sequence Change Counts are defined per DSID and hence are only meaningful in the context of a single DSID.

The CMTS MUST assign Packet Sequence Numbers to packets from the same Service Flow being transmitted using the same DSID in the order that these packets were classified to the Service Flow. The CMTS MUST increment Packet Sequence Numbers by 1 for each packet transmitted using the DSID. All sequenced packets transmitted with the same DSID on a particular downstream channel MUST be transmitted by the CMTS with strictly increasing Packet Sequence Numbers. The CMTS MUST transmit sequenced packets only on channels included in the Downstream Resequencing Channel List for the DSID.

Due to differences in internal CMTS transmission latency for different downstream channels, the CMTS may initially transmit sequenced packets on a set of downstream channels already slightly out of order. The CMTS SHOULD start transmission to a downstream interface of sequenced packets with the same DSID with no more than the "Default" CMTS Skew between an earlier higher packet sequence number and a later lower packet sequence number. The CMTS MUST start transmission to a downstream interface of sequenced packets with the same DSID with no more than the "Maximum" CMTS Skew between an earlier higher packet sequence number and a later lower packet sequence number. Default and Maximum CMTS Skew are defined in annex B.

The CM MUST forward packets from the resequencing operation for further processing in order of increasing Packet Sequence Number.

For a particular DSID, the CM's Next Expected Packet Sequence Number is defined as the sequence number which is one greater than the Packet Sequence Number of the last packet forwarded for further processing. For a newly created Resequencing DSID without associated multicast encodings, the CM MUST initialize its Next Expected Packet Sequence Number to zero. When the CM first begins receiving a Resequencing DSID with associated multicast encodings or in the event of a change in Sequence Change Count (clause 8.2.3) on a DSID the CM is already receiving, the CM MUST choose an initial value for Next Expected Packet Sequence Number based on actual received Packet Sequence Numbers. The algorithm for choosing this initial value is vendor-specific. When choosing an initial value for Next Expected Packet Sequence Number, the CM MAY discard otherwise valid packets and/or delay forwarding of packets on the DSID for the duration of Max_Resequencing_Wait from the time it first begins receiving packets on the DSID (in the case of a new DSID) or from the time it first receives a packet with the new Sequence Change Count (in the case of a change in Sequence Change Count). If the CM discards packets when choosing an initial value for Next Expected Packet Sequence Number, it MUST NOT generate CM-STATUS messages or increment any MIB error counters based on these discards. Certain Resequencing DSIDs might be created during Registration specifically for a single CM, yet contain multicast encodings for use with individually-directed multicast packets (see clause 9.2.2.5). Although the CMTS does not explicitly indicate to the CM that such a DSID has been created exclusively for the CM, the first packet labeled with this type of DSIDs will be given sequence number zero (0). In order to provide reliable service (particularly for eSAFE provisioning traffic), the CM SHOULD minimize any packet loss when choosing an initial value for Next Expected Packet Sequence Number for DSIDs that are communicated to the CM during Registration.

The CM MUST define an internal, vendor-specific Resequencing Window Size which is less than or equal to 2^{15} for a given DSID. This value is determined by the CM vendor and is not signaled by the CMTS. The Resequencing Window Size has units of packets and approximately represents the number of packets the CM is able to simultaneously store for resequencing on a particular DSID. The vendor may define this parameter based on various device-specific characteristics such as maximum throughput supported, number of downstream channels supported, etc. For example, for a device which supports P packets per second on each downstream channel and has D downstream channels, the Resequencing Window Size could be chosen as $(P * \text{Max_Resequencing_Wait} * D)$. Max_Resequencing_Wait refers to the maximum value of DSID Resequencing Wait Time as described in annex B.

The CM MUST store a received DSID-labeled packet with a Packet Sequence Number which is greater than or equal to the Next Expected Packet Sequence Number for the DSID and less than or equal to the Next Expected Packet Sequence Number plus the Resequencing Window Size for the DSID. Such a Packet Sequence Number is defined to be "in-range".

If the Packet Sequence Number of a received in-range DSID-labeled packet is equal to the Next Expected Packet Sequence Number, the CM SHOULD immediately forward it for further processing and increment the Next Expected Packet Sequence Number by 1. If the Next Expected Packet Sequence Number now matches the Packet Sequence Number of another stored packet, the CM SHOULD immediately forward this packet for further processing as well and again increment its Next Expected Packet Sequence Number. This process repeats until the CM's Next Expected Packet Sequence Number does not match the Packet Sequence Number of any currently stored packet.

If the Packet Sequence Number of a received in-range DSID-labeled packet is not equal to the Next Expected Packet Sequence Number, the CM determines that some sequence numbers are "missing". Missing sequence numbers are those which are less than the Packet Sequence Number of the packet just received, greater than or equal to the Next Expected Packet Sequence Number and not already received and stored by the CM. The CM MUST wait at least the DSID Resequencing Wait Time for a missing sequence number to arrive. This interval begins at the time of completion of arrival of the packet which first caused the missing packet to be identified as missing. The CM is allowed to wait longer than the DSID Resequencing Wait Time, but it SHOULD minimize the amount of time it waits beyond the specified value. If a packet is received and the CM waited longer than the Resequencing Warning Threshold but less than the DSID Resequencing Wait Time, the CM increments a Resequencing Warning Counter.

If the CM waits the required interval for a missing sequence number and the missing sequence number does not arrive, the CM declares the missing sequence number to be "lost".

When the Next Expected Packet Sequence Number is declared lost, the CM MUST perform the following sequence of actions:

- 1) Increment the Next Expected Packet Sequence Number until it is not a number which has been declared lost.
- 2) If the new value of Next Expected Packet Sequence Number matches the Packet Sequence Number of a currently stored packet, forward this packet for further processing and return to step 1, otherwise end.

The CM associates a Downstream Resequencing Channel List with each Resequencing DSID. This may be explicitly signaled in a Downstream Resequencing Channel List subtype encoding of the Resequencing Encoding of a DSID. If it is not explicitly signaled, it is set equal to the Receive Channel Set of the CM. Per clause 9.1.2.2, the CM will drop a DSID-labeled packet arriving on a downstream channel which is not part of the Downstream Resequencing Channel List associated with that DSID.

Whenever the CM has stored a sequenced packet on all active channels of the Downstream Resequencing Channel List of a Resequencing DSID, the CM declares all sequence numbers lower than the lowest stored sequence number to be lost. This is termed "rapid loss detection". When packets are declared lost in this manner, the CM MUST set its Next Expected Packet Sequence Number equal to the lowest stored sequence number. The CM MUST then forward stored packets in order and increment the Next Expected Packet Sequence number accordingly until the Next Expected Packet Sequence Number does not match the sequence number of a currently stored packet. The CMTS MAY transmit a "Sequenced Null Packet" (clause 6.2.5.6, DS-EHDR) on an otherwise idle downstream channel to facilitate rapid loss detection.

8.2.3.2 Skew Requirements

In downstream channel bonding, there are multiple physical paths between the CMTS and a given CM. These paths may have different delays. This delay variation results in "skew" across the CM's received channels.

For purposes of this clause, each possible path from the CMTS bonding distribution point to the CM's RF input is modeled as consisting of four components:

- 1) CMTS internal queuing/processing delays.
- 2) CIN delay and EQAM internal queuing/processing delays.
- 3) Downstream interleaver delay.
- 4) Physical delays (e.g. propagation delay, group delay) on the HFC plant itself.

Of these four components, only items 1 through 3 may vary significantly across channels and/or from one packet to the next; hence, only these items contribute to skew. Only the physical HFC plant delay from the CMTS or EQAMs to a given CM may be considered fixed (i.e. any variations are on the order of microseconds and are small compared to the total skew).

The following skew budget was used to arrive at the skew time values in this specification:

- Variation in CMTS internal queuing/processing delays (CMTS Skew): 3,0 ms.
- Variation in CIN delay and EQAM internal delay: 4,5 ms.
- Variation across channels in downstream interleaver delay: 10,5 ms (i.e. the difference in delay between (128,4) and (32,4) for J.83 annex B QAM 256 downstream channels).

In a particular deployed system, a different skew budget may apply for any of a number of reasons (e.g. type of CMTS (integrated vs. modular), downstream interleaver settings, etc.). A different skew budget may also apply for traffic carried within different DEPI Packet Streaming Protocol (PSP) flows, since the delay through a particular CIN may be different for each flow. The operator is free to use a different skew budget as long as the skew seen at the CM's RF input is within the CM's capabilities as specified below. If the skew seen at the CM's RF input is not within the CM's capabilities, proper system operation cannot be expected.

Because of skew, a packet transmitted by the CMTS with a lower Packet Sequence Number may arrive at the CM later than a packet with a higher Packet Sequence Number. Such packets are called "out of order" sequenced packets. The difference between the arrival times of these packets at the output of the CM's deinterleaver is termed "CM Skew". CM Skew is defined to be the difference in the completion of arrival of all symbols of out-of-order sequenced packets at the Downstream RF input interface of the CM, plus the difference in the end-to-end delay of the downstream interleaver on different downstream channels.

Due to differences in internal CMTS transmission latency for different downstream channels, the CMTS may initially transmit bonded packets on a set of downstream channels already slightly out of order. "CMTS Skew" is defined as the interval between the start of transmission of out-of-order sequenced packets as measured at the set of CMTS [4] and [2] interfaces.

The DSID Resequencing Wait Time is a per-DSID signaled value from the CMTS to the CM (see clause C.1.3.1.31). It indicates how long a CM will wait for "missing" out-of-order packets to arrive. Its use is detailed in clause 8.2.3.1. The CMTS selects the DSID Resequencing Wait Time for each DSID based on the expected maximum value of the CM skew for the DSID. Each DSID may have a different DSID Resequencing Wait Time due to differing downstream channels in the various bonded channel sets, as well as differing CIN delays from different DEPI flows. When the Resequencing Channel List for the DSID changes, it is possible that the DSID Resequencing Wait Time will change as well. The CMTS may use the DOCSIS Path Verify (clause 10.5.1) mechanism as a tool for determining an appropriate DSID Resequencing Wait Time value. The DSID Resequencing Wait Time value may change over time, e.g. due to changes in loading in the CIN, reconfiguration of the CIN or other changes in plant conditions. The CMTS may discover these changes based on DPV measurements or as a result of provisioning changes by the operator. The CMTS MUST select a value for DSID Resequencing Wait Time that is within the range specified in TLVs in annex C.

NOTE: A larger DSID Resequencing Wait Time may translate into increased latency at the CM and reduced system performance. Hence, it is desirable to keep skew to a minimum. In an M-CMTS, the operator should ensure that packets from any given service flow receive similar QoS treatment in the CIN, especially if these packets are sent on different DEPI flows. This will minimize the skew contribution of the CIN.

8.2.3.3 Resequencing DSID Signalling

The Downstream Resequencing Encoding of a DSID Encoding (clause C.1.5.4.3) defines the following attributes for a DSID:

- Resequencing Enabled.
- Downstream Resequencing Channel List.
- DSID Resequencing Wait Time.
- Resequencing Warning Threshold.
- CM-STATUS holdoff timer for out-of-range events.

The Resequencing Enabled subtype indicates whether the DSID requires a resequencing context in the CM. The Downstream Resequencing Channel List provides a list of Downstream Channel IDs on which the CM resequencing context performs rapid loss detection. The DSID Resequencing Wait Time is used by the CM to determine when packets are "lost" as described in clause 8.2.3.1. The Resequencing Warning Threshold is used as a threshold for counting and reporting. The CM-STATUS Maximum Holdoff Timer parameter controls the reporting of packets with out-of-range sequence numbers as described in clause 6.4.34.

The CMTS MUST receive confirmation (via REG-ACK or DBC-RSP) that a CM has added the DSID before transmitting packets labeled with a Resequencing DSID that does not have associated multicast subtype encodings.

8.2.4 Cable Modem Physical Receive Channel Configuration

A Cable Modem reports its ability to receive multiple channels with one or more Receive Channel Profile (RCP) Encodings in a REG-REQ or REG-REQ-MP message. Each Receive Channel Profile describes a logical representation of the CM's downstream physical layer in terms of Receive Channels (RCs) and Receive Modules (RMs). The CMTS initially configures a CM's Receive Channels and Receive Modules with a Receive Channel Configuration (RCC) Encoding in the REG-RSP or REG-RSP-MP Message. This clause defines the applicable terms and outlines the mechanism by which this process takes place.

8.2.4.1 Receive Channels

The term "Receive Channel" refers to the component of a Cable Modem that receives a single Downstream Channel on a single center frequency. A CM is considered to implement a fixed number of Receive Channels, each of which is identified within the CM by a Receive Channel Identifier. The CMTS assigns one or more of its Downstream Channels to the Receive Channels of a CM by assigning the center frequency of the Receive Channel in a Receive Channel Configuration Encoding. The CMTS MUST assign the Receive Channels of a CM to the Downstream Channels which are in a single MAC Domain.

A Receive Channel Profile communicated from CM to CMTS defines the following attributes of each Receive Channel:

- Index, a 1-based index that identifies the Receive Channel (required).
- Connection Capability, a bit map that provides the set of one or more higher level Receive Modules to which the Receive Channel can connect.
- Connected Offset, for the case when the Receive Channel connects to a single Receive Module (e.g. a demodulator group) that defines a block of adjacent channels, this attribute defines the 1-based offset of the Receive Channel within that block.
- Primary Downstream Channel Capability, a boolean that indicates whether the Receive Channel is capable of providing the DOCSIS master clock reference to the CM.
- Vendor Specific Capabilities (optional).

A Receive Channel Configuration communicated from CMTS to CM assigns the following attributes to a Receive Channel:

- Center Frequency Assignment, the center frequency defining the single DOCSIS downstream channel for the Receive Channel (required).
- Primary Downstream Channel Indicator, a boolean that indicates that the CMTS assigns the Receive Channel to provide master clock reference timing to the CM.
- Connection Assignment, the single member from the set of higher level Receive Modules described in a Connection Capability RCP encoding to which the CMTS assigns the Receive Channel to actually connect.
- Vendor Specific Configuration (optional).

8.2.4.2 Receive Modules

The term "Receive Module" refers to a component in the CM physical layer implementation shared by multiple Receive Channels. Examples of Receive Modules include analog tuners, intermediate frequency down-converters, analog-to-digital converters, digital sample buses and digital signal processing modules. A Receive Module in a Receive Channel Profile represents the constraints on channel assignment caused by the common component. The purpose for identifying the Receive Modules in a Receive Channel Profile is to communicate those constraints to the CMTS and to permit the CMTS to reconfigure the frequencies of Receive Channels while minimizing the disruption of data received by the CM.

Whenever the CM is forced to reconfigure a shared physical layer component during normal operation, a disruption may occur on all receive channels sharing that component. The reconfiguration may cause a data error on any packets being received through the shared component. For example, a reconfiguration to a shared component serving the CM's Primary Downstream Channel may cause the CM to lose DOCSIS master clock synchronization, possibly forcing re-ranging on the upstream channels.

A goal of DOCSIS downstream channel bonding is to permit the CMTS to rapidly change the assignment of a CM's receive channels (e.g. for load balancing or IP television channel changes) with minimal packet loss. In some cases the CMTS can change a CM's Receive Channel Set without forcing the CM to reconfigure a shared physical component.

A shared physical component causes a dependency on the group of receive channels sharing that component. For example, an analog tuner component forces all receive channels sharing that tuner to have center frequencies within the range of the tuner.

A Receive Channel is said to "connect" to a Receive Module when it uses the shared component. Depending on the CM implementation and the type of physical component, the connection from a Receive Channel to a Receive Module is either fixed or configurable. If the connection is fixed, the CM communicates the fixed connection in the RCP. In this case, the Connection Capability attribute of the Receive Channel indicates a single Receive Module. If the connection is configurable, the CM communicates in the RCP the set of multiple Receive Modules to which a Receive Channel is capable of connecting. In this case, the Connection Capability attribute of the Receive Channel indicates more than one Receive Module. When a connection is configurable, the CMTS in the RCC assigns the Receive Channel to connect to one particular Receive Module. A Receive Module may have no Receive Channels connected to it.

The following are examples of a Receive Module in a CM:

- A limited capture bandwidth analog tuner.
- An A/D converter for an adjacent band of channels.
- A multiple-channel digital signal processing block.
- A single CM chip within a subscriber device that contains multiple CM chips.

The first three examples require the CMTS to assign the set of Receive Channels to a limited range of center frequencies. The last example requires the CMTS to limit downstream channel bonding to only the Receive Channels of the same CM chip.

A Receive Channel Profile communicated from CM to CMTS defines one or more of the following capability attributes of a Receive Module:

- Index, a 1-based index to identify the Receive Module (required).
- Number of Adjacent Channels, when the Receive Module describes a component that serves a block of adjacent channels, e.g. for an analog tuner or a demodulator group, this attribute defines the number of such adjacent channels.
- Channel Block Range, when the adjacent channel block for the Receive Module described above is limited to a subset of the full DOCSIS frequency range (e.g. for an analog tuner), this configuration provides the minimum center frequency of the first channel in the block and the maximum center frequency of the last channel in the block.
- Resequencing Capable Subset, the set of Receive Channels that may be defined in the same Resequencing Channel List of a DSID for downstream channel bonding.
- Common Physical Layer Parameters, the set of physical layer parameters such as modulation type or interleaver that are shared by all Receive Channels connected to the Receive Module.
- Connection Capability, a bit map that provides the set of one or more higher level Receive Modules to which this Receive Module can connect.
- Vendor-specific capabilities (optional).

A Receive Channel Configuration communicated from CMTS to CM assigns one or more of the following attributes to a Receive Module:

- First Channel Center Frequency, for a Receive Module defining a block of adjacent channels, this parameter assigns the center frequency of the lowest-frequency channel of the block.
- Connection Assignment, the single member from the set of higher level Receive Modules described in a Connection Capability RCP encoding to which the CMTS assigns the Receive Channel to actually connect.
- Vendor Specific Configuration, corresponding to vendor-specific capabilities.

A CMTS is expected, but not necessarily required, to assign Receive Channels connected to a Receive Module that defines a block of adjacent channels to center frequencies located at an integral number of channel widths from the first channel center frequency of the block.

8.2.4.2.1 Receive Module Interconnection

Some CM architectures may support the concept of a programmable interconnection between a Receive Channel and a Receive Module. For example, a Receive Channel may be programmed to be connected to only one of several different A/D converters. Furthermore, a Receive Module itself (e.g. an A/D converter) may be programmed to be connected to one of several different "higher level" Receive Modules (e.g. one of a set of analog tuners with fixed frequency ranges). In other cases, a Receive Channel will have a fixed interconnection to a Receive Module (e.g. the third channel of a digital signal processing component encompassing four adjacent channels).

Receive Channels connect to Receive Modules and Receive Modules can connect to an arbitrary number of "higher level" Receive Modules (i.e. Receive Modules closer to the RF interface connector). The CM **MUST NOT** report a Receive Channel Profile with a loop of capable Receive Module connections.

In a Receive Channel Configuration, the CMTS configures Receive Channels to frequencies and assigns the connections between Receive Channels and Receive Modules such that all of the constraints of the Receive Module are met.

Figure 8-1 depicts the interconnection between Receive Channels and Receive Modules in a Receive Channel Profile.

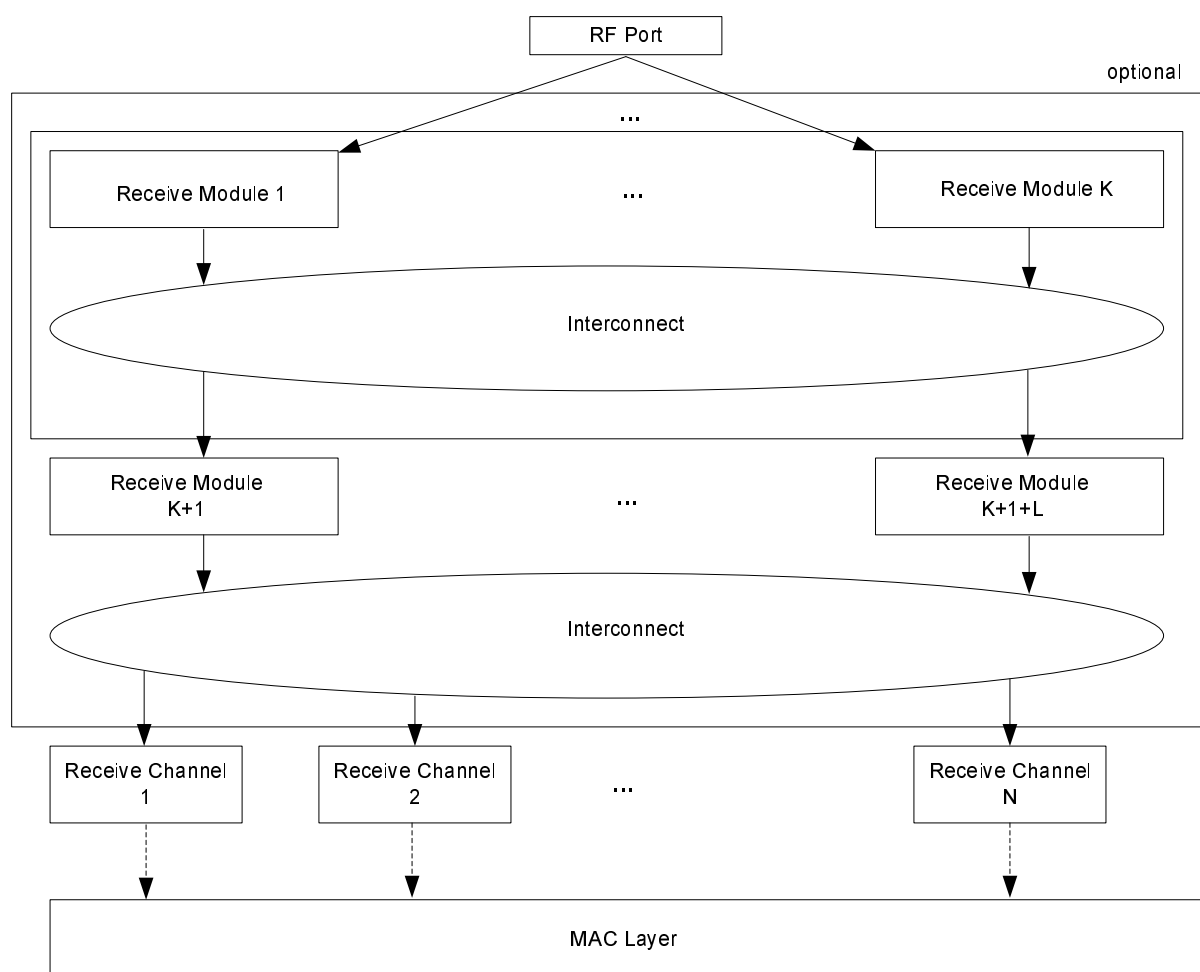


Figure 8-1: Interconnection between Receive Channels and Receive Modules

In a Receive Channel Profile, a Receive Channel is considered to be a physical layer component at the "lowest" level (i.e. the farthest from the RF connector). Each Receive Channel delivers a sequence of DOCSIS MAC frames from a single center frequency. A Receive Channel Profile describes a fixed number of Receive Channels, numbered consecutively from 1.

In a Receive Channel Profile, any layers of Receive Modules above the Receive Channels are optional. A multiple-channel CM may be implemented with Receive Channels that have no dependencies on other channels. Such a CM would not describe any Receive Modules and the Receive Channels would be considered to connect directly to the physical RF port of the CM.

When Receive Modules are present, the interconnection between Receive Channels and Receive Modules may either be fixed by the CM implementation or configurable by the CMTS. If the interconnection is configurable, the particular Receive Module to which an individual Receive Channel is connected may affect the dependency between the channels in the CM. For example, if a Receive Channel can be configured to one of multiple Receive Modules, the choice of a particular Receive Module could limit the set of frequencies to which the Receive Channel can be moved without disrupting other channels.

8.2.4.3 Receive Channel Profile

A Receive Channel Profile is an encoding that represents the Receive Channels and Receive Modules (if any) of the CM. A CM **MUST** communicate to the CMTS one or more Receive Channel Profile (RCP) Encodings within its Registration Request, using the TLV structure as defined in clause C.1.5.3.

A Receive Channel Profile is defined for operation with either 6 MHz or 8 MHz center frequency spacing. The CMTS advertises in its periodic MAC Domain Descriptor (MDD) messages a Receive Channel Profile Reporting Control TLV clause 6.4.28.1.4, that controls how the CM reports RCPs in its REG-REQ message. One subtype of this TLV is the RCP Center Frequency subtype that controls whether the CM should report RCPs based on 6 MHz or 8 MHz center frequency spacing. When the CM registers with the CMTS, it sends only the Receive Channel Profiles defined for the requested spacing. The CM **MUST** communicate to the CMTS all of the Standard Receive Channel Profiles (see annex E) that are defined for the requested spacing and that are supported by the CM.

A Receive Channel Profile is identified with a globally unique five-byte Receive Channel Profile Identifier (RCP-ID) consisting of two parts:

- the 3-byte Organization Unique Identifier (OUI) assigned to the CM manufacturer by the IEEE; and
- a 2-byte Manufacturer Receive Channel Profile Identifier assigned by the CM manufacturer uniquely to the profile.

The CMTS advertises in its MDD message that contains a Verbose RCP Reporting subtype clause 6.4.28, MAC Domain Descriptor, to request that the CM report a verbose description of the RCPs. The verbose description contains complete sub-type encodings to describe Receive Channels and Receive Modules. If a verbose description is not requested, the CM reports only the Receive Channel Profile Identifiers.

In order to reduce cable operator configuration requirements, a CM **MAY** report a manufacturer-specific RCP-ID using the 3-byte OUI and 2-byte RCP Profile Identifier assigned by a CM silicon manufacturer.

If a CM advertising "DOCSIS 3.0" in the DOCSIS Version Capability (clause C.1.3.1.2) does not communicate an RCP to the CMTS, the CMTS **MUST** reject the registration.

8.2.4.3.1 Standard Receive Channel Profiles

In order to avoid requiring CMTS software to support an increasing number of arbitrarily complex RCPs, DOCSIS defines the concept of a Standard RCPs. A Standard RCP represents a well-known virtual model of Receive Channels and Receive Modules that describes a useful minimum feature set for a class of multiple-channel DOCSIS subscriber devices.

The Standard RCPs defined by an organization are assigned identifiers with the organization's OUI. See annex E for the definition of the set of DOCSIS Standard RCPs at the time of release of this specification. New Standard RCPs may be defined at any time, independent of the revision process of the present document. Refer to the CableLabs web site for a list of Standard RCPs defined by CableLabs. Other organizations may define additional Standard RCPs.

For example, CLAB-6M-004A describes four Receive Channels of 6 MHz width assigned by the CM to a single Receive Module that restricts the assignment of the Receive Channels to fall within a 60 MHz bandwidth (i.e. a range of 10 adjacent channels). This is depicted in figure 8-2, which follows.




Figure 8-2: Standard Receive Channel Profile CLAB-6M-004A

A CM can support multiple RCPs, including Manufacturer RCPs and Standard RCPs. The CM **SHOULD** include all supported RCPs (with the relevant center frequency spacing) in its Registration Request to the CMTS.

8.2.4.4 Receive Channel Configuration

The CMTS **MUST** select one of the RCPs in the Registration Request for configuring the CM. The CMTS returns, in a Registration Response to a CM, a "Receive Channel Configuration" (RCC) Encoding that contains TLVs to configure the Receive Channels and Receive Modules of the selected RCP. clause C.1.5.3 describes the TLVs of an RCC. The RCC provides the particular RCP-ID that the CMTS selected for configuring the CM.

For example, the RCC for the Standard RCP CLAB-6M-004A will configure:

- the center frequency of Receive Channels 1 through 4 (modulation, annex and interleaver depth are optional); and
- the Starting Center Frequency of Receive Module 1, i.e. where the 10-channel adjacent channel group is placed in the downstream spectrum.

A CMTS is not required to select a Manufacturer RCP for the RCC. The CMTS is permitted to always select a Standard RCP for configuration. If the CM reports an RCP that is supported by the CMTS, the CMTS **MUST** send an RCC encoding in the Registration Response. If the CM does not send a Verbose RCP and the CMTS does not recognize any of the RCP-IDs advertised by the CM, the CMTS **MUST NOT** send an RCC in the REG-RSP or REG-RSP-MP. If the CM does not receive an RCC encoding in the Registration Response, it **MUST** set its Receive Channel Set to only contain its current Primary Downstream Channel.

When autonomous load balancing is disabled for a CM, the CMTS is required to assign the CM an RCC in which the Primary Downstream Channel matches the CM's candidate Primary Downstream Channel as defined in clause 11.6.2. When autonomous load balancing is enabled for a CM, a valid RCC need not contain the CM's candidate Primary Downstream Channel. The CMTS **MUST NOT** send an RCC containing any Receive Channel which is in a different MAC Domain than the CM's candidate Primary Downstream Channel.

An "active" Receive Channel is defined to be one configured with a Receive Channel Center Frequency encoding in an RCC. The CMTS **MAY** omit a Receive Channel encoding for a Receive Channel Index, in which case the CM does not activate the Receive Channel. The CMTS **MUST NOT** transmit an invalid RCC encoding. A valid RCC is one that meets the following requirements:

- It contains exactly one RCP-ID.
- Any Receive Module configuration (clause C.1.5.3.4) is for a Receive Module Index defined for the selected RCP.
- Any Receive Channel configuration (clause C.1.5.3.5) is for a Receive Channel Index defined for the selected RCP.

- Any Receive Module First Channel Center Frequency Assignment (clause C.1.5.3.4.4) defines a frequency within the minimum and maximum range of center frequencies configured for any Receive Module to which the Receive Channel connects.
- A Receive Channel Connectivity Assignment Encoding (clause C.1.5.3.5.2) exists in the RCC for each Receive Channel Connectivity Capability encoding in the RCP when the Receive Channel is configured as active.
- A Receive Module Connectivity Assignment Encoding (clause C.1.5.3.4.6) exists in the RCC for each Receive Module Connectivity Capability encoding in the RCP when the RM connects directly or indirectly (via other RMs) from any active Receive Channel.
- Any Receive Module Connectivity Assignment Encoding (clause C.1.5.3.4.6) in the RCC connects a Receive Module to exactly one of the choices described in the Receive Module Connectivity Capability encoding of the RCP.
- Any Receive Channel Connectivity Assignment Encoding (clause C.1.5.3.5.2) in the RCC connects a Receive Module to exactly one of the choices described in the Receive Channel Connectivity Capability encoding of the RCP.
- A Receive Module First Channel Center Frequency Assignment (clause C.1.5.3) exists for a Receive Module that reports an Adjacent Channel capability and is connected to an active Receive Channel.
- Any Receive Channel Center Frequency Encoding (clause C.1.5.3) matches any Receive Channel Connected Offset for an active Receive Channel connected to a Receive Module with Adjacent Channels (clause C.1.5.3).
- Any Receive Channel Center Frequency Encoding is within the range defined by DOCSIS and on a channel configured for a Downstream Channel on the CMTS.
- When the CMTS knows the MAC Domain Downstream Service Group (MD-DS-SG) for a CM, any Receive Channel Center Frequency Encoding communicated to that CM corresponds to a Downstream Channel configured to reach that MD-DS-SG.
- Exactly one Receive Channel is configured with the Receive Channel Primary Downstream Channel Indicator (clause C.1.5.3.5.5) enabled.
- The physical layer parameters of all downstream channels assigned to Receive Channels connected to the same Receive Module match any Receive Module Common Physical Layer Parameter encoding in the RCP for that Receive Module.

If an RCC is invalid, the CM rejects the REG-RSP, REG-RSP-MP or DBC-REQ message that contains the invalid RCC.

8.2.4.4.1 Static Receive Module Assignments

The placement of Receive Modules in the downstream spectrum and the interconnection between Receive Channels and Receive Modules can require arbitrary complexity in the CMTS. To avoid this, the CMTS MAY support the static configuration of the parameters and interconnections of a Receive Module.

[10] defines the objects for configuring static Receive Module assignments.

A CMTS MAY limit RCC assignments to only the Receive Modules statically configured by the Cable Operator. For example, a CMTS may require a Cable Operator to statically configure the starting center frequency of the Receive Modules for all RCPs of interest.

A static Receive Module assignment may not assign all Receive Module parameters. For example, it may assign the interconnections between Receive Channels and Receive Modules, but not assign the first Receive Channel Frequency of a Receive Module.

A Cable Operator may configure multiple static Receive Modules for the same RCP-ID. In this case, the CMTS selects any one of the relevant static Receive Modules.

8.2.4.5 RCC Message Sequence Example

Figure 8-3 depicts an abbreviated version of the CM Initialization sequence showing the communication of CM Receive Channel information in a DOCSIS system.

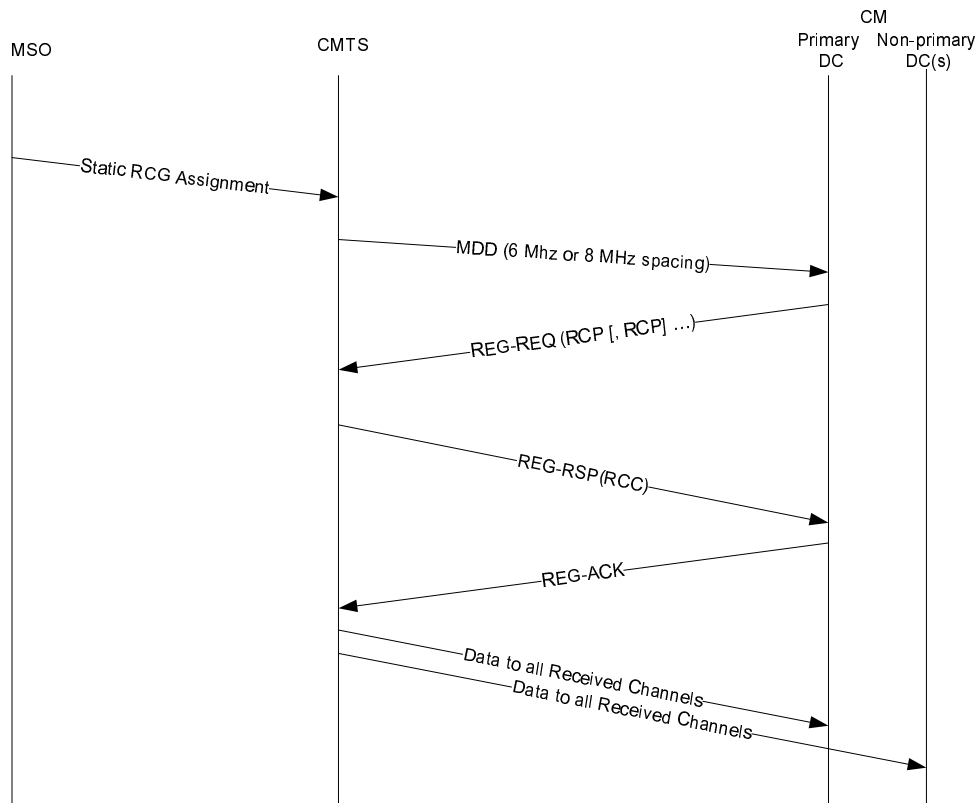


Figure 8-3: Receive Channel Configuration Message Sequence

- 1) The Cable Operator configures a CMTS with a set of static Receive Module (RM) assignments, e.g. assigning the First Center Frequency of the 10-channel wide Receive Module of Standard RCP CLAB-6M-004A.
- 2) The CMTS regularly broadcasts a MAC Domain Descriptor message that indicates that the CM should report Receive Channel Profiles for a particular center frequency spacing (i.e. 6 MHz or 8 MHz).
- 3) The CM sends a Registration Request message to the CMTS containing its Receive Channel Profile(s) (RCP) at the indicated center frequency spacing.
- 4) The CMTS selects one of the CM's RCPs and returns a Registration Response message with a Receive Channel Configuration (RCC) Encoding for that RCP, selected from the set of RCCs that are appropriate for this RCP.
- 5) The CM configures its Receive Channels according to the RCC.
- 6) When the CM has completed processing of the Registration Response, it returns a Registration Acknowledgement message.
- 7) When the CMTS receives the Registration Acknowledgement message, it begins transmitting data to the Receive Channels of the CM.

8.2.5 QoS for Downstream Channel Bonding

While the CMTS is required to maintain the DOCSIS Quality of Service for a Bonded Service Flow, the actual output data burst size for a Bonded Service Flow at the CMCI port may differ from the Maximum Traffic Burst QoS parameter for the flow. This is a result of CMTS packet distribution process, the CM resequencing operation and (in the case of an M-CMTS), variable delays in the CIN. The CM is required to wait for late arriving packets and once the CM completes resequencing a set of received packets (either by receiving the next expected packet or by expiry of its resequencing timer) it may release the set of packets in a single burst, see clause C.2.2.5.3.

8.3 Upstream Channel Bonding

An upstream bonding group consists of two or more upstream channels over which a service flow may be transmitted. A service flow may be assigned to a single upstream channel or an upstream bonding group.

Multiple Transmit Channel Mode (MTC Mode) provides mechanisms and capabilities that enable Upstream Channel Bonding. If a CM is operating in MTC Mode, all of its service flows, whether assigned to a single channel or to an upstream bonding group, operate with the mechanisms that are supported in MTC Mode (clause C.1.3.1.24). Compared to pre-3.0 DOCSIS operation, request mechanisms, grant mechanisms and grant-filling mechanisms are different for MTC Mode operation. In MTC Mode, CMs make queue-depth based requests for a service flow and the CMTS decides how to allocate grants to that service flow over the upstream channels usable for that service flow. Request mechanisms are described in clause 7.2.1.4.

8.3.1 Granting Bandwidth

The CMTS scheduler allocates bandwidth on the individual channels based on the available bandwidth on all of the bonded upstream channels. Requests transmitted on any individual channel may be allocated bandwidth on any combination of upstream channels within the bonding group associated with the requesting service flow. In this manner, the CMTS can perform real-time load balancing of the upstream channels. Similarly, the CMTS can consider the physical layer parameters on each of the upstream channels and the requested number of bytes to figure out the optimal allocations across channels.

8.3.2 Upstream Transmissions with Upstream Channel Bonding

For upstream channel bonding, the CM uses segmentation with Continuous Concatenation and Fragmentation (CCF) to fill the grants allocated to each service flow. The CM **MUST NOT** combine different service flows within a segment. CCF uses a segment header to aid the CMTS in reconstructing the original data sent for each service flow. For some unsolicited grant services the CM does not need to fragment, so a segment header is not needed to aid in reassembly for these services. In order to reduce the overhead for these services, the use of segment headers is enabled or disabled on a per service flow by using the Request/Transmission Policy.

Regardless of whether Segment Headers are enabled or disabled for a service flow, the CMTS **MAY** allocate SIDs for more than one upstream channel in the SID cluster associated with the service flow. Regardless of whether Segment Headers are enabled or disabled for a service flow, the CM **MUST** be prepared to transmit on any upstream channel for which a SID has been allocated by the CMTS in the SID cluster.

8.3.2.1 Segment Header ON Operation

Each service flow for Multiple Transmit Channel Mode operation is provisioned for either Segment Header ON operation or Segment Header OFF operation. With Segment Header On operation, the CM **MUST** place the 8-byte segment header at the beginning of every segment for the service flow. For the first segment transmitted on a given Service Flow after that flow is configured with a non-null AdmittedQosParamSet, the CM **MUST** set the Sequence Number field of the Segment Header to zero. The segment header format is defined in clause 6.3.

When the CM makes a bandwidth request, it **MUST NOT** include the segment header overhead in its request, since the CM has no idea how many grants the CMTS may use (and thus how many segment headers to assume) in granting the request. The CMTS **MUST** account for the segment overhead when granting requests to service flows provisioned for Segment Header ON operation.

8.3.2.2 Segment Header OFF Operation

For service flows with Segment Header OFF, the CM MUST NOT use the fragmentation portion of CCF. For service flows with Segment Header OFF, the CM MUST NOT use the concatenation portion of CCF. Thus, all segments transmitted by the CM for these service flows MUST contain only a single complete packet. If a segment is lost, the CMTS MAC will know that the next segment boundary aligns with a packet boundary and can continue processing the received packets for that service flow.

In the absence of explicit, vendor-specific configuration to the contrary, the CMTS MUST NOT allocate bandwidth on more than one upstream channel for a given Segment Header OFF service flow. The reason for this restriction is that the packet ordering across channels cannot generally be guaranteed without segment headers.

Note that segment-header-off operation is permitted only for unsolicited grant services. Unsolicited grant services can be configured for either Segment Header ON or Segment Header OFF operation.

If a CMTS receives a Registration Request message with a Service Flow configured with Segment Header OFF from a CM that will be operating in Multiple Transmit Channel Mode, the CMTS MUST reject the Registration Request if the service flow is neither UGS nor UGS-AD. If a CMTS receives a DSA Request message with a Service Flow configured with Segment Header OFF for a CM that is operating in Multiple Transmit Channel Mode, the CMTS MUST reject the DSA Request if the service flow is neither UGS nor UGS-AD.

For a Service Flow with Segment Header Off, piggyback requesting is not allowed since the scheduling type is either UGS or UGS-AD.

8.3.3 Dynamic Range Window

The Dynamic Range Window defines a 12 dB range of Transmit Power Levels the CM can use for each of the channels in its Transmit Channel Set. The DRW is communicated from the CMTS to the CM in the RNG-RSP or in the TCC encodings of the REG-RSP-MP or the DBC-REQ message. The top of the DRW is defined as $P_{load_min_set}$ [12] and is expressed as some number of dB below P_{hi} for each channel. Because P_{hi} for each channel depends on the Modulation Technology (TDMA or S-CDMA) and Modulation Order for the channel as well as the total number of active channels, some changes in the burst profiles for a channel or a change in the number of active channels could cause the Tx Level (P_t) that the modem was using to lie outside the Dynamic Range Window. Scenarios in which that could happen and associated CM requirements are described below. The CMTS is required to manage the Dynamic Range Window for the modem and if the CMTS commands the modem to do something which would result in a violation of the Dynamic Range Window, the CM will reject or ignore the command.

8.3.3.1 Channels Added During Registration

The CMTS is required to send a value for the Dynamic Range Window in the Registration Response message if the CM is to be operating in Multiple Transmit Channel Mode. The CM MUST use the Dynamic Range Window value sent in the Registration Response. When the CM receives the REG-RSP-MP it determines which upstream channels it will be using based on the TCC encodings. The CM then collects UCDs for the designated channels and determines P_{hi} for each of the channels. The CM determines what transmit power level it would be using on each channel in the new TCS after applying any power offsets commanded by the TCC. If the Dynamic Range Window value communicated to the CM in the TCC would cause the transmit level for any of the channels in the commanded Transmit Channel Set to lie outside the Dynamic Range Window, the CM MUST re-initialize the MAC with an Initialization Reason of DYNAMIC-RANGE-WINDOW-VIOLATION (19).

If the CM is able to use the TCC encodings and if it needs to adjust the Dynamic Range Window, it will wait for a Global Reconfiguration Time [12] prior to beginning the ranging process (or sending the REG-ACK if the Initialization Technique for all the channels is "Use Directly").

If Power Offset TLVs were provided in the TCC encodings the following rules will apply:

- If the Initialization Technique for any channel requires ranging the CM MUST begin ranging using the Tx Level determined by applying the commanded offset.
- If the Initialization Technique is "Use Directly" for any channels in the TCS, the CM MUST use the Tx Levels determined by applying the commanded offsets. If a Global Reconfiguration Time is needed in order to apply a commanded Tx Level, the CM will wait for a Global Reconfiguration Time [12] before using the channel.

If no Power Offset TLVs were provided, the CM will begin ranging using Power Offset values stored in non-volatile memory if values exist for the channels and if those values lie within the Dynamic Range Window. On those channels for which no Power Offset TLV is provided and no valid value is stored in non-volatile memory, the CM will set its Tx Level at the bottom of the Dynamic Range Window and begin ranging with that value. If the modem undergoes T3 timeouts during initial ranging it will adjust its Tx Levels in a vendor specific manner and attempt to range using other levels within the Dynamic Range Window, leaving no power level range greater than 6 dB untried until it receives a RNG-RSP, clause 10.2.3.7.1.

The CM MUST NOT at any time set its Tx Level to a value that would lie outside the Dynamic Range Window. If the modem is able to use some, but not all, of the channels and at least one of those channels is a channel that is associated with the Primary US Service Flow, the CM registers using partial-service. In the event that the CM is unable to acquire some of the channels and goes into partial-service, the CM MUST maintain the P_{hi} values that it calculated based on the number of channels commanded in the TCC encodings and not recalculate P_{hi} based on the number of channels it is actually able to use. If none of the channels associated with the primary US Service Flow in the TCS are useable, the CM MUST re-initialize the MAC with an Initialization Reason of NO_PRIM_SF_USCHAN (17).

8.3.3.2 Channels Added by a DBC-REQ

If a Dynamic Range Window value is provided in the DBC-REQ, the CM MUST use that value. If no Dynamic Range Window value is provided in the DBC-REQ, the CM continues to use the value that it had been using. When the CM receives the DBC-REQ it determines which upstream channels it will be using based on the TCC encodings. The CM then collects UCDs for the designated channels and determines P_{hi} for each of the channels. The CM determines what transmit power level it would be using on each channel in the new TCS after applying any power offsets commanded by the TCC. If the Dynamic Range Window value communicated to the CM in the DBC-REQ would cause the Tx Level for any of the channels in the commanded Transmit Channel Set to lie outside the Dynamic Range Window, the CM MUST send a DBC-RSP with a Confirmation Code of reject-dynamic-range-window-violation (210) and continue operation.

If the CM is able to use the TCC encodings and if it needs to adjust the Dynamic Range Window, it will wait for a Global Reconfiguration Time [12] prior to beginning the ranging process if ranging is required for any of the channels.

If Power Offset TLVs were provided in the TCC encodings the following rules will apply:

- If the Initialization Technique for any channel requires ranging, the CM MUST begin ranging using the Tx Level determined by applying the commanded offset.
- If the Initialization Technique is "Use Directly" for any channels being added to the TCS, the CM MUST use the Tx Levels determined by applying the commanded offsets. The CM will begin using those channels immediately unless the Tx Level for a particular channel lies outside the current Dynamic Range Window, in which case it will wait for a Global Reconfiguration Time [12] before using the affected channel.
- If the Initialization Technique is "Use Directly" for any channels the CM is already using, the CM will continue uninterrupted use of the channels, meaning that any Tx Level adjustments resulting from applying the Power Offsets would be handled the same as adjustments provided in a RNG-RSP.

If no Power Offset TLVs were provided, the CM will begin ranging using Power Offset values stored in non-volatile memory if values exist for the channels and if those values lie within the Dynamic Range Window. On those channels for which no Power Offset TLV is provided and no valid value is stored in non-volatile memory, the CM will set its Tx Level at the bottom of the Dynamic Range Window and begin ranging with that value. If the modem undergoes T3 timeouts during initial ranging it will adjust its Tx Levels in a vendor specific manner and attempt to range using other Tx Levels within the Dynamic Range Window, leaving no power level range greater than 6dB untried until it receives a RNG-RSP, clause 10.2.3.7.1.

The CM MUST NOT at any time set its Tx Level to a value that would lie outside the Dynamic Range Window. If the modem is able to use some, but not all, of the channels and at least one of those channels is a channel that is associated with the Primary US Service Flow, the CM enters partial-service. In the event that the CM is unable to acquire some of the channels and goes into partial-service, the CM MUST maintain the P_{hi} values that it calculated based on the number of channels commanded in the TCC encodings and not recalculate P_{hi} based on the number of channels it is actually able to use. If none of the channels associated with the primary US Service Flow in the TCS are useable, the CM MUST re-initialize the MAC with an Initialization Reason of NO_PRIM_SF_USCHAN (17).

8.3.3.3 Channels Deleted by a DBC-REQ

A DBC-REQ that deletes channels from the CM's Transmit Channel Set could result in an increase in P_{hi} for the remaining channels. When the CM receives a DBC-REQ that deletes some of the upstream channels, the modem recalculates P_{hi} based on the number of remaining channels and the UCDs for those channels. If the Dynamic Range Window value communicated to the CM in the DBC-REQ would cause the Tx Level for any of the channels in the commanded Transmit Channel Set to lie outside the Dynamic Range Window, the CM MUST send a DBC-RSP with a Confirmation Code of reject-dynamic-range-window-violation (210) and continue operation with the unmodified Transmit Channel Set.

8.3.3.4 UCD Changes Burst Profiles Resulting in New Value for P_{hi}

If the CM receives a new UCD with burst profile changes such that P_{hi} for the channel is changed, the CM MUST adjust its Reported Transmit Power (P_r) for the channel by an amount equal to the change in P_{hi} such that P_{load} [12] is maintained. By definition, this adjustment in P_r will result in the CM maintaining the same delta with respect to the top of the Dynamic Range Window as the CM was using prior to the UCD change.

8.3.3.5 Power Offset in RNG-RSP Causing Dynamic Range Window Violation

If the CM receives a RNG-RSP with a Power Level Adjust TLV or a Power Offset TLV that would cause a violation of the Dynamic Range Window, the CM MUST ignore the commanded adjustment and indicate an error condition (clause 6.4.5).

8.4 Partial Service

Whenever one or more channels in the Transmit Channel Set (TCS) and/or the Receive Channel Set (RCS) are unusable, that CM is said to be operating in a "partial service" mode of operation in the upstream and/or downstream respectively. A channel is deemed to be unusable when the CM is unable to acquire one or more channels during registration and/or DBC or if a CM lost an upstream and/or downstream channel during normal operation. It is intended to be a temporary mode of operation where services may not operate normally and which can be resolved via several means.

The CM signals that it is in a partial service mode of operation to the CMTS via the appropriate means:

- The REG-ACK if the channel is not acquired during registration.
- The DBC-RSP if the channel is not acquired during Dynamic Bonding Change.
- The CM-STATUS message, if a channel becomes unusable during normal operation.

When an upstream channel is unusable, the CM MUST NOT use any request, data or broadcast initial maintenance opportunities. The CM MUST respond to any unicast ranging opportunities on an unusable upstream channel in order to attempt to establish or re-establish communications on that channel. The CM is no longer in a partial service mode of operation in the upstream when there are no unusable upstream channels. This occurs when the CM receives a RNG-RSP(success) for all of the channels in the TCS or unusable upstream channels are removed from the TCS via DBC messaging such that the CM is no longer operating with a subset of its TCS.

When a non-primary downstream channel is unusable, the CM MUST continue to attempt to acquire those downstream channels. Note that if the CM is unable to acquire the primary downstream channel during registration or DBC, the CM will immediately perform a MAC re-init. Also note that if the CM loses the primary downstream channel during normal operations, it will cease transmitting on all upstream channels, but will continue to attempt to re-acquire the primary downstream channel until another timeout (such as for periodic ranging) causes a re-init MAC operation. The CM is no longer in a partial service mode of operation in the downstream when there are no unusable downstream channels. This occurs when the CM is able to acquire or re-acquire all of the channels in the RCS or unusable downstream channels are removed from the RCS via DBC messaging such that the CM is no longer operating with a subset of its RCS.

When the CMTS is aware that an upstream channel is unusable, the CMTS **MUST NOT** provide unicast transmission opportunities for the CM other than ranging opportunities for that upstream channel. Likewise, when the CMTS is notified by the CM that a downstream channel is unusable, the CMTS **MUST NOT** transmit unicast packets destined for that CM or its CPEs on that downstream channel. When the CM is operating on only a subset of its TCS and/or RCS, the CMTS **SHOULD** attempt to meet minimum QoS guarantees and maintain poll/grant intervals, but is not required to do so. The CMTS **SHOULD** attempt to resolve partial service situations, such as by providing the CM opportunities to acquire or re-acquire the affected channels or via DBC messaging.

9 Data Forwarding

This clause defines the rules and requirements for CM and CMTS forwarding in the DOCSIS 3.0 Network. There are primarily 3 types of packets that a DOCSIS network is concerned with forwarding: broadcast (IPv4 packets destined for all hosts), multicast (IPv4 or IPv6 packets sent to a group of hosts) or unicast (IPv4 or IPv6 packets destined for a single host). An IPv6 anycast address is syntactically indistinguishable from a unicast address. Therefore, throughout the rest of the present document references to unicast addresses also apply to anycast addresses. This specification has been limited to focus on IPv4 and IPv6 network layer protocols. Other protocols could be supported, but their operation is not specified.

The DOCSIS 3.0 CMTS uses the DSID (see clause 7.4) as a labeling technique to differentiate certain traffic types and to prevent modems and hosts from receiving packets that they are not intended to receive. The CMTS communicates the appropriate DSID label to each CM. In some instances, the CM uses the DSID to forward packets destined for the CM and any devices behind the CM.

9.1 General Forwarding Requirements

The data-over-cable system transmits Internet Protocol version 4 and/or version 6 (IPv4 and/or IPv6) packets transparently between the head-end and the subscriber location.

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 uses 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. A CMTS using link layer forwarding is known as a bridging CMTS. A CMTS using network layer forwarding is known as a routing CMTS.

Data forwarding through the CM is link-layer transparent bridging. Forwarding rules are similar to [16] with modifications to allow for the support of multiple network layers. Provisions exist in this specification for frames to be passed from a higher-layer entity (such as the SNMP agent or DHCP client within the CM) to be forwarded by the cable modem.

CMs **MAY** support the [16] spanning tree protocol of [20] with the modifications described in annex K. The CM **MUST** include the ability to filter (and disregard) [16] Bridge Protocol Data Units (BPDUs). The CMTS **MUST** include the ability to filter (and disregard) [16] BPDUs.

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).
- TFTP (Trivial File Transfer Protocol), which is used by the modem for downloading operational software and configuration information.
- DHCP (Dynamic Host Configuration Protocol) v4 and v6, frameworks for passing configuration information to hosts on a TCP/IP network.
- HTTP (HyperText Transfer Protocol), which is optionally used by the modem for downloading operational software.

Certain management functions also use IP. These management functions include, for example, supporting spectrum management.

The protocol stacks at the CM and CMTS RF interfaces are shown in figure 9-1.

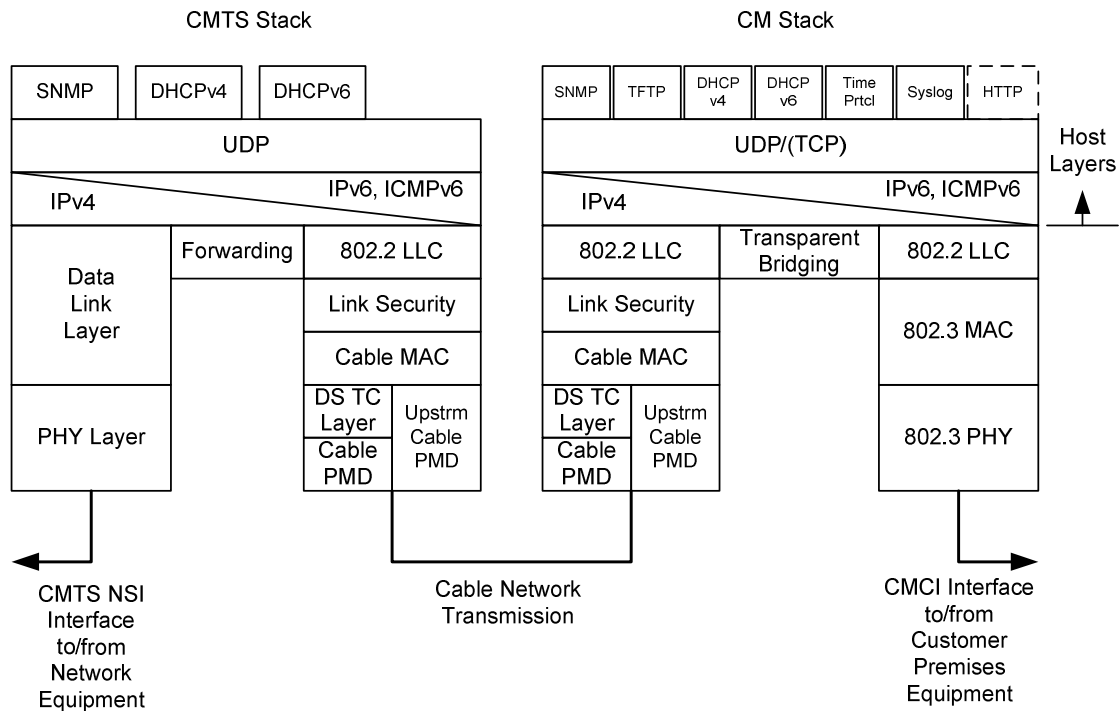


Figure 9-1: DOCSIS Protocol Stacks

9.1.1 CMTS Forwarding Rules

9.1.1.1 General CMTS Forwarding

Data forwarding through the CMTS **MUST** be transparent bridging, network-layer forwarding (routing, IP switching) or a combination of the two. The CMTS **MUST** provide IP (v4 and v6) connectivity between hosts attached to cable modems and do so in a way that meets the expectations of Ethernet-attached customer equipment. For IPv6, the CMTS is not required to deliver traffic between hosts attached to different cable modems using link-local scope addresses.

The CMTS **SHOULD** replicate broadcast packets on all primary-capable Downstream Channels of a MAC Domain. A CMTS may provide a proxy ARP service to avoid forwarding ARP (see [15]) messages. A proxy ARP service on the CMTS reduces the possibility of potential denial of service attacks because the ARP messages are not forwarded to hosts (untrusted entities). The implementation of the proxy ARP service is vendor dependent.

For IPv6 the CMTS **SHOULD** either forward Neighbor Discovery (ND) packets [42] on primary-capable Downstream Channels of the MAC domain or facilitate ND-based services (also known as "proxy ND service") to avoid forwarding ND messages. A proxy ND service on the CMTS reduces the possibility of potential denial of service attacks because the ND messages are not forwarded to hosts (untrusted entities). The implementation of the proxy ND service is vendor dependent.

Because the CMTS is not required to track MLD messages forwarded by CMs that are not MDF-enabled, the CMTS may have incomplete knowledge of solicited node multicast addresses in use on the CMTS RFI at any time. For example, an initializing CM could send two MLD membership reports for Solicited Node Multicast Groups prior to being considered MDF-enabled by the CMTS. Additionally, MDF-disabled CMs or MDF-incapable CMs may indicate support for IPv6 and as such may operate in IPv6 provisioning mode and/or may support IPv6 eSAFes/CPEs. When the CMTS needs to transmit a packet addressed to a solicited node multicast address and if the CMTS does not know which primary downstream(s) to use, the CMTS **MUST** transmit the packet on every primary capable downstream that is in the link-local scope of the packet.

A CMTS that supports routing of IPv6 traffic is not required to support advertisement of not on-link ([59]) prefix assignment, which eliminates the use of ND for non-link-local scope address resolution.

If the CMTS transparently bridges data, the CMTS MUST pad out the PDU and recompute the CRC for PDUs less than 64 bytes to be forwarded from the upstream RFI. The CMTS and CM MAY support the forwarding of other network layer protocols other than IP. If the forwarding of other network layer protocols is supported, the ability to restrict the network layer to IPv4 and IPv6 MUST be supported by the CMTS.

At the CMTS, if link-layer forwarding is used, then it MUST conform to the following general [16] rules:

- The CMTS MUST NOT duplicate link-layer frames.
- The CMTS MUST deliver link-layer frames on a given Service Flow, clause 6.1.2.3, in the order they are received subject to the skew requirements in clause 8.2.3.2.

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 [34] with respect to its CMTS-RFI and CMTS-NSI interfaces.

A bridging CMTS applies appropriate DSID labeling and forwarding of the packets received from the NSI interface according to the rules in clause 9.1.1.2, DSID labeling and pre-registration multicast. The NSI-side router generates the Router Advertisement (RA) message [16] to the RFI interface for appropriate DSID labeling and forwarding by the bridging CMTS.

A bridging CMTS MUST forward the packets destined to the well-known IPv6 MAC addresses (see annex A) to the NSI-side router for processing.

A routing CMTS applies appropriate DSID labeling and forwarding of the packets received from the NSI interface according to the rules in clause 9.1.1.2, DSID labeling and pre-registration multicast. When the routing CMTS forwards well-known IPv6 multicast packets from the NSI to RFI, the CMTS terminates and applies appropriate processing for these packets. The routing CMTS generates the RA message [42] for appropriate DSID labeling and forwarding to the RF interface.

The CMTS MUST discard IPv6 RA messages received on its RF interface.

9.1.1.2 DSID Labeling

In addition to its forwarding responsibilities, the CMTS labels packets it forwards to the CM with a DSID according to the following rules:

- The CMTS SHOULD NOT label broadcast packets (addressed to a MAC destination of FF:FF:FF:FF:FF:FF) with a DSID.
- The CMTS labels multicast packets according to the rules specified in clause 9.2.2.2.
- The CMTS MAY label traffic bearing an individual MAC destination address with a DSID to indicate its resequencing context. The CMTS SHOULD NOT label traffic bearing an individual MAC destination address with a DSID if that traffic is not sequenced.

However, in cases such as virtual private networks, the above rules need not apply and the CMTS MAY label traffic with a DSID to limit the interpretation of layer 2 MAC addresses to a "virtual LAN" of CMs on the RF MAC interface.

9.1.2 CM Address Acquisition, Filtering and Forwarding Rules

The CM MUST support forwarding of IP traffic (both IPv4 and IPv6). CMs and CMTSs operate as IP and LLC hosts as defined by [16] for communication over the cable network.

The term "CPE MAC addresses" used in this clause includes MAC addresses of both connected CPE devices and eSAFEs. The term "CMCI port" describes physical interfaces to which connected CPE devices can attach. The term "Logical CPE Interface" refers to an interface between the CM and an eSAFE. The term "CPE port" refers to an interface that is either a CMCI port or a Logical CPE Interface.

Data forwarding through the CM is link-layer bridging with the rules specified in the following clauses.

9.1.2.1 MAC Address Acquisition

The CM maintains a forwarding database (bridging table) including entries for the CM's own MAC address and CPE MAC addresses.

The CM MUST acquire CPE Ethernet MAC addresses, either from the provisioning process or from learning, until the CM acquires its maximum number of CPE MAC addresses (the lesser of the Max CPE from the config file (clause C.1.1.7, Max CPE) or 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 64 CPE MAC addresses.

The CM MUST NOT learn any MAC addresses for its forwarding database prior to registration. The CM MUST allow configuration of CPE MAC addresses during the provisioning process (up to its maximum number of CPE addresses) to support configurations in which learning is not practical, nor desired. The CM MUST give provisioned addresses precedence over learned addresses when adding entries to the forwarding database. The CM MUST NOT age out CPE MAC addresses. The CM MUST place all acquired CPE MAC addresses in its forwarding database [32].

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), the CM MUST discard all provisioned and learned addresses.

9.1.2.2 CM Filtering Rules

The CM MUST discard frames that are received with CRC or frame format errors. The CM MUST discard packets based on the configurable filtering mechanisms defined in [9] and clause 7.5.1.2.2.

Filtering downstream frames received on any of the downstream channels in the CM's Receive Channel Set conforms to the following specific rules:

- The CM MUST discard frames with an unknown SAID.
- The CM MUST discard unicast frames addressed to unknown destination MAC addresses (MAC addresses not contained in the CM's forwarding database), even if the SAID is known. The CM MUST NOT generate a TEK Invalid (see [15]) or report a CRC error in this case.
- If Multicast DSID Forwarding is enabled (reference clause C.1.3.1.33), the CM MUST discard all packets (unicast, multicast and broadcast) with a DS EHDR containing an unknown DSID value (even if the MAC destination address or SAID is known). The CM MUST NOT generate a TEK Invalid (see [15]) due to a key sequence error or report a CRC error in this case. Additional CM requirements for the forwarding of unicast, multicast and broadcast packets that apply when MDF is disabled are detailed in annex G.
- The CM MUST discard all DSID labeled packets which are labeled with a Resequencing DSID and received on a downstream channel not in the Downstream Resequencing Channel List associated with the DSID.
- The CM MUST discard multicast frames from source addresses which are provisioned or learned as supported CPE devices.
- The CM MUST discard broadcast frames from source addresses which are provisioned or learned as supported CPE device.
- The CM MUST discard broadcast frames not labeled with a DSID which are received on any channel other than the CM's Primary Downstream Channel.

Forwarding of frames received from any CPE port to the RFI conforms to the following specific rules:

- The CM MUST NOT transmit upstream frames from source MAC addresses other than those provisioned or learned as supported CPE devices.
- The CM MUST NOT transmit upstream Router Advertisements (RAs) received on any interface.

9.1.2.3 CM Forwarding Rules

The CM MUST NOT duplicate link-layer frames.

9.1.2.3.1 CM Pre-Operational Forwarding Behavior

Prior to becoming operational as in figure 10-1, the CM operates per the following rules:

- The CM MUST forward to its IP stack all unicast frames that are received on the Primary Downstream Channel and addressed to the CM's MAC address.
- The CM MUST forward from its IP stack to the RF interface the multicast traffic that is necessary for completing the registration process.
- The CM MUST NOT send any DHCPv4 DHCPDISCOVER or DHCPREQUEST, DHCPv6 Solicit or Request, TFTP-RRQ, HTTP Request, Time Protocol Request or IPv6 Router Solicitation messages to any interface except the RF Interface.
- The CM MUST NOT accept any DHCPv4 DHCPOFFER or DHCPACK, DHCPv6 Advertise or Reply, TFTP-DATA, HTTP Response, Time Protocol Response or IPv6 Router Advertisements from the CMCI ports.
- The CM MUST NOT forward any packets from the RF interface to any CPE port.
- The CM MUST NOT forward any packets from any CPE port to the RF Interface.

9.1.2.3.2 CM Operational Forwarding Behavior

Once the CM is operational as in figure 10-1, CM forwarding in the upstream and downstream directions conforms to the following rules:

- The CM MAY perform one or more frame/packet processing functions on frames received from the CPE port prior to classifying them to a Service Flow. Example frame/packet processing functions include: DOCSIS protocol filtering as specified in [10], a policy-based filtering service as described in clause 7.5.6.1 and annex L and priority-based queuing to support 802.1P/Q services. Unless specified otherwise, the CM MUST transmit upstream link-layer frames in the order that they are received on a given Service Flow. The CM SHOULD support a mechanism by which TCP ACK frames are prioritized or filtered in order to increase TCP session throughput.
- Unless specified otherwise, the CM MUST deliver downstream sequenced link-layer frames for a particular DSID in the order indicated by the Packet Sequence Number (see clause 8.2.3.1). The CM MUST deliver downstream non-sequenced link-layer frames of the same traffic priority in the order that they are received on a given downstream channel. Relative packet ordering of such frames received on different downstream channels is not specified (see clause 8.2.1).
- The CM MAY perform one or more frame/packet processing functions on frames received from the RF port prior to transmitting them on the CPE port. Example frame/packet processing functions include: DOCSIS protocol filtering as specified in [10], a policy-based filtering service as described in clause 7.5.6.1 and annex L and priority-based queuing to support 802.1P/Q services.
- The CM MUST NOT forward frames between the RF port and CPE ports if the CM config file sets Network Access Control Object (NACO) to 0. The CM MUST forward frames between the CPE ports and CM IP stack even if NACO is 0. The CM MUST forward frames between the RF port and CM IP stack even if NACO is 0.

Forwarding of non-DSID labeled downstream frames received on any of the downstream channels in the CM's Receive Channel Set conforms to the following specific rules:

- The CM MUST forward unicast frames addressed to the CM's MAC address to the CM's IP stack.
- The CM MUST forward unicast frames addressed to learned MAC addresses to the CPE port on which the address was learned.
- The CM MUST forward unicast frames addressed to provisioned MAC addresses to all CPE ports, until that MAC address is learned on a particular CPE port.

- The CM MUST forward broadcast frames not labeled with a DSID which are received on the Primary Downstream Channel to the CPE ports and the CM IP stack.

Forwarding of DSID-labeled downstream frames received on any of the downstream channels in the CM's Receive Channel Set conforms to the following specific rules:

- The CM MUST forward unicast packets which are labeled with a known DSID and addressed to the CM's MAC address to the CM's IP stack.
- The CM MUST forward unicast packets labeled with a known DSID to the CPE port on which the destination MAC address was learned.
- The CM MUST forward unicast frames which are labeled with a known DSID and addressed to provisioned MAC addresses to all CPE ports, until that MAC address is learned on a particular CPE port.
- A CM MUST forward broadcast packets labeled with a known DSID to only the union of: all interfaces identified in the Multicast CM Interface Mask associated with that DSID; and all interfaces identified by the list of client MAC addresses associated with that DSID.

Forwarding of frames received from any CPE port conforms to the following specific rules:

- The CM MUST forward frames addressed to unknown destination MAC addresses only to the RF Interface.
- The CM MUST forward broadcast frames to all ports (including the CM IP stack) except the port which received the frame.
- The CM MUST forward frames addressed to known destination MAC addresses to the port on which the destination address was learned.
- The CM MUST NOT accept any DHCPv4 DHCPOFFER or DHCPACK, DHCPv6 Advertise or Reply, TFTP-DATA, HTTP Response, Time Protocol Response or IPv6 Router Advertisements from any of the CPE ports for the purposes of configuration, secure software download or address renewal.

Forwarding of frames received from any CMCI port(s) conforms to the following specific rules:

- The CM MUST forward multicast frames to the RF port, the CM IP stack and all CMCI ports except the port which received the frame.
- The CM MUST NOT forward multicast frames to any Logical CPE Interfaces.

Forwarding of frames received from any Logical CPE Interface conforms to the following specific rules:

- The CM MUST forward multicast frames to the RF port.
- The CM MUST NOT forward multicast frames to any ports other than the RF port.

Forwarding of frames being sent by the CM IP stack conforms to the following specific rules:

- The CM MUST forward frames addressed to unknown destinations only to the RF port.
- The CM MUST forward broadcast frames to all ports.
- The CM MUST forward multicast frames to the RF port.
- The CM MUST NOT forward multicast frames to any ports other than the RF port.
- The CM MUST forward frames to the port on which the destination address was learned.
- The CM MUST NOT forward any DHCPv4 DHCPDISCOVER or DHCPREQUEST, DHCPv6 Solicit or Request, TFTP-RRQ, HTTP Request, Time Protocol Request or Router Solicitation messages to any ports except the RF port.

9.2 Multicast Forwarding

9.2.1 Introduction

Multicast can provide significant bandwidth savings in a network. Multicast is especially attractive in the cable network because of the broadcast nature of the cable downstream. In addition to providing end to end bandwidth savings, the cable RF network can be used effectively to distribute multicast streams to multiple downstream devices. With the introduction of channel bonding in DOCSIS 3.0 the potential scope of multicast applications in the cable network is much greater than with earlier DOCSIS implementations.

DOCSIS 3.0 defines a flexible infrastructure for multicast that can accommodate a wide range of new protocols and services. For example, this specification supports both the traditional form of IP Multicast referred to as "Any Source Multicast" (ASM) (as defined in [29]), as well as "Source Specific Multicast" (SSM). SSM is particularly relevant for broadcast-type IP multicast applications as it offers additional security due to the single source nature of SSM. IGMPv3 [52] and MLDv2 [54] are required for SSM. In addition, there is a potential to leverage this infrastructure in conjunction with technologies such as PacketCable Multimedia [23] for offering new applications or services. This infrastructure can also be used to offer Layer 2 Virtual Private Networking [8] services.

DOCSIS 1.1/2.0 relied on the snooping of IGMPv2 messaging by the CM. By snooping in the CM, the ability to move to newer multicast technologies was limited. In order to enable the flexibility and scalability to support a large array of multicast protocols, DOCSIS 3.0 defines the cable modem to be multicast protocol agnostic and introduces centralized control at the CMTS. This approach simplifies the cable modem operation and reduces the overall cost of deploying multicast solutions. However, in order to ensure that a DOCSIS 3.0 cable modem can operate in a Pre-3.0 DOCSIS environment, the CM is still required to snoop IGMPv2 messages when operating with a Pre-3.0 DOCSIS CMTS.

The Multicast Model, shown in figure 9-2 contains various entities that control the multicast subsystem at the CMTS such as IGMP and MLD for dynamic operation and configuration through CLI or SNMP for static operation. Other entities may include PIM [56] and 802.1Q, GARP/GMRP [16]. These entities can trigger the CMTS to signal a DSID along with a set of group forwarding attributes to specific CMs based on events such as IGMP joins.

A CMTS-initiated control mechanism replaces the IGMPv2 snooping and the associated multicast filtering in the cable modem in earlier DOCSIS versions, as indicated by the control path in figure 9-2. From the CMTS perspective, a DSID identifies a subset of CMs intended to receive the same Multicast session. From the CM perspective, the DSID is a filtering and forwarding criterion for multicast packets. The group forwarding attributes associated with a DSID enable or disable the forwarding of multicast packets to specific interfaces in the cable modem.

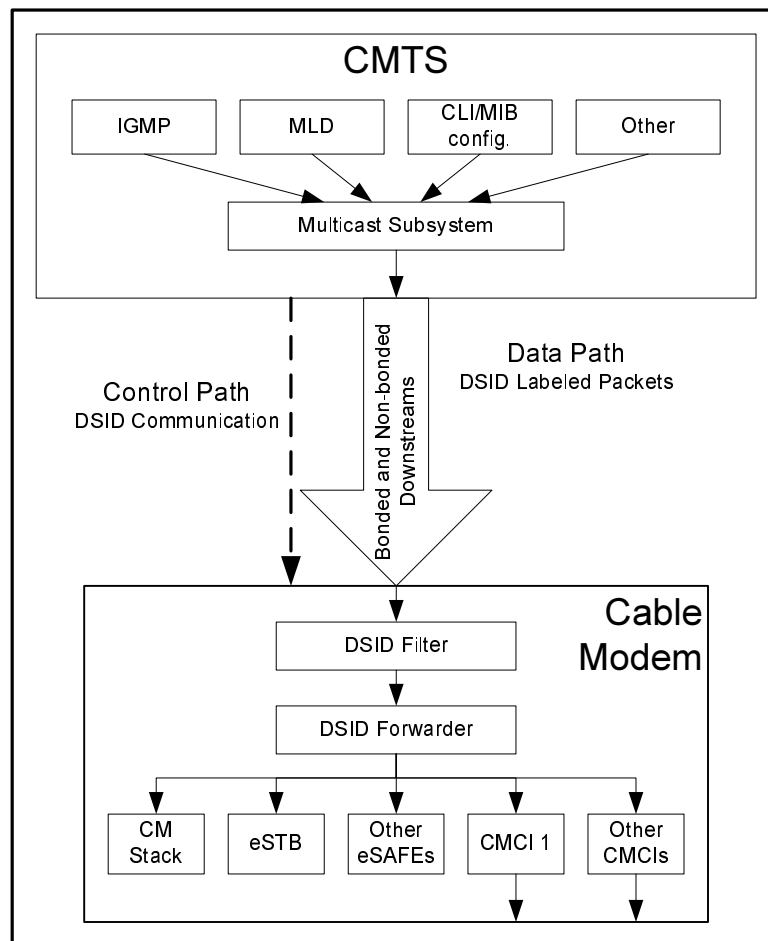


Figure 9-2: Multicast Model

9.2.2 Downstream Multicast Forwarding

This clause outlines the CMTS requirements when the Multicast DSID Forwarding is enabled on the CMTS. This clause also outlines the CM requirements when the CMTS sets the Multicast DSID Forwarding Capability of a CM to GMAC-Promiscuous(2).

Annex G identifies exceptions or enhancements to the CM requirements described in this clause when the CMTS sets the Multicast DSID Forwarding capability of a CM to either GMAC-Explicit(1) or Disabled(0). Annex G also identifies CMTS requirements when Multicast DSID Forwarding is disabled on the CMTS.

A CMTS is said to enable Multicast DSID Forwarding on a MAC Domain when it enables Multicast DSID forwarding to any CM registered on that MAC domain. A CMTS is said to disable MDF forwarding on a MAC Domain when it disables Multicast DSID Forwarding to all CMs registered on that MAC Domain. A CMTS that returns a non-zero value of the Multicast DSID Forwarding Support capability encoding to a CM in a REG-RSP or REG-RSP-MP is said to "enable" Multicast DSID Forwarding at the CM. Although a CM reports that it is capable of Multicast DSID Forwarding, the CMTS may return a value of 0 for the encoding in its REG-RSP or REG-RSP-MP. The CMTS is said to "disable" DSID Multicast Forwarding in this case.

The CMTS considers a CM to be "MDF-capable" when the CM reports a non-zero value for the capability of "Multicast DSID Forwarding" in REG-REQ or REG-REQ-MP. The CMTS considers a CM to be "MDF-incapable" when the CM reports a zero value for the capability of "Multicast DSID Forwarding" in REG-REQ or REG-REQ-MP.

An MDF-capable CM is considered to operate in one of the following three modes of operation based on the value set by the CMTS in REG-RSP or REG-RSP-MP for the Multicast DSID Forwarding (MDF) Capability, see clause C.1.3.1.33:

- When the CMTS sets the value of 0 for MDF capability, the CM is considered to operate in "MDF-disabled Mode". The CM and CMTS requirements for this mode of operation are detailed in clause G.4.3.

- When the CMTS confirms the value of 1 for MDF capability, the CM is considered to operate in "GMAC-Explicit MDF Mode". The CM and CMTS requirements for this mode of operation are detailed in clause G.4.2.
- When the CMTS sets or confirms the value of 2 for MDF capability, the CM is considered to operate in "GMAC-Promiscuous MDF Mode". GMAC-Promiscuous MDF Mode means that the CM has the ability to "promiscuously" accept and forward all GMAC addresses with known DSID labels. 3.0 CMs and later are required to implement and advertise the capability of MDF=2. The requirements for both CM and CMTS for GMAC-Promiscuous MDF Mode are detailed in the following clauses.

There are two main classes of IP multicast traffic that need to be forwarded by the DOCSIS 3.0 CMTS: traffic associated with the well known IPv6 groups (annex A) when IPv6 forwarding is configured and user-joined multicast. User-joined multicast is defined as multicast traffic that is based on IGMP or MLD protocols where clients and routers have defined messages that are used to start and stop the reception of multicast traffic.

Downstream multicast packet forwarding at the CM is achieved by filtering and forwarding packets based on DSIDs. This involves the following three high level functions:

- 1) Labeling multicast packets with a DSID by the CMTS.
- 2) Communicating DSIDs and associated group forwarding attributes to a CM by the CMTS.
- 3) Filtering and forwarding of DSID labeled multicast packets by the CM.

The term "IP Multicast Session" is used to refer to both ASM IP multicast groups and SSM IP multicast channels. The term JoinMulticastSession is used to refer to an IGMP/MLD message element that indicates a "join to an ASM IP multicast group" or a "subscribe to an SSM IP multicast channel". The term LeaveMulticastSession is used to refer to an IGMP/MLD message element that indicates a "leave from an ASM IP multicast group" or an "unsubscribe from an SSM IP multicast channel". The term "Multicast Client" refers to an entity with a unique MAC address that receives multicast packets (e.g. CM IP Host stack, e-SAFE devices or CPE devices connected to the CM).

9.2.2.1 Examples of Downstream Multicast Forwarding using DSIDs

DOCSIS 3.0 introduced the capability of CMs to receive multiple Downstream Channels (DCs) and therefore to receive multicast session traffic distributed by a CMTS on a Downstream Bonding Group (DBG) of multiple channels. CMs incapable of receiving multiple Downstream Channels can receive multicast traffic on only a single Downstream Channel. Because DOCSIS 3.0 supports MAC domains of multiple downstream channels with a mixture of both single-receive-channel and multiple-receive-channel CMs, it poses the special problem of avoiding the duplicate delivery of downstream multicast traffic. For example, when a multicast session is replicated to separate downstream channels in order to reach DOCSIS 2.0 CMs on each channel, a DOCSIS 3.0 CM that receives both channels needs to avoid delivering both copies of the packet to its CPE interface.

An important concept with Multicast DSID-based Forwarding is the Downstream Channel Set (DCS). A Downstream Channel Set is defined as either: a single Downstream Channel (DC) or a Downstream Bonding Group (DBG) of more than one channel. Each Downstream Channel Set is composed of downstream channels in a single MAC Domain. With DOCSIS 3.0, the CMTS forwards IP Multicast packets received on a Network System Interface (NSI) to one or more Downstream Channel Sets of a CMTS MAC Domain.

For purposes of downstream DSID-based Multicast Forwarding, a "bonding CM" is considered to be one that has a non-zero Multiple Receive Channel Support capability set by the CMTS as described in clause C.1.3.1.29. A "nonbonding CM" is considered to be one that has the Multiple Receive Channel Support capability set to zero by the CMTS.

Multicast DSID-based Forwarding avoids undesired duplicate delivery of IP multicast session traffic by using the DSID label to distinguish each replication of an IP multicast session to a particular set of CMs.

The example in figure 9-3 depicts the use of DSIDs to prevent duplicate delivery of two non-bonded multicast sessions by a bonding CM to its CPE(s).

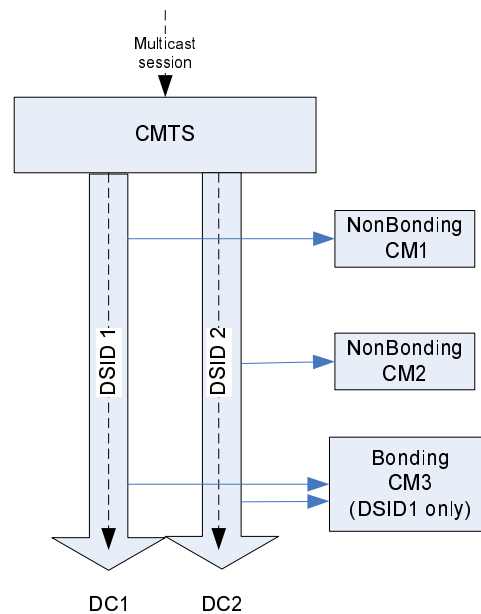


Figure 9-3: DSIDs prevent duplication of non-bonded replications

Figure 9-3 depicts a CMTS that receives a multicast session and replicates it on downstream channel DC1 to reach nonbonding CM1 and replicates it to downstream channel DC2 to reach nonbonding CM2. The Ethernet Packet PDUs transmitted on each downstream channel are identical, i.e. they have the same layer 2 Group MAC destination address and the same layer 3 IP contents. The only difference is in the Downstream Service Extended Header (DS-EHDR) that the CMTS prepends on the MAC frames on each channel. The CMTS labels the DS-EHDR of the replicated frames on DC1 with DSID1 and labels the DS-EHDR of the replicated frames on DC2 with DSID2. The nonbonding CMs ignore the DSID label and forward the replication received on their (single) Primary Downstream Channel. The CMTS instructs bonding CM3, however, to forward multicast traffic labeled only with DSID1 and does not inform CM3 of the value of DSID2 at all. CM3 therefore forwards the replicated traffic on DC1 (and labeled with DSID1) and discards the replicated traffic on DC2 because it is labeled with the unknown label DSID2.

The CMTS uses DSIDs in a similar way to restrict forwarding of source-specific multicast sessions through only the CMs with multicast clients that have joined the SSM session. An SSM session is identified by the pair (S,G) for a multicast source S sending to an IP multicast group G. Because DOCSIS 1.1/2.0 CMs filter downstream multicast traffic based only on the destination group G, they forward multicast traffic for both (S1,G) and (S2,G) to their CPE ports. CMs capable of Multicast DSID-based Forwarding (MDF), however, can use DSID filtering to limit forwarding to a single (S,G) session. This is depicted in figure 9-4.

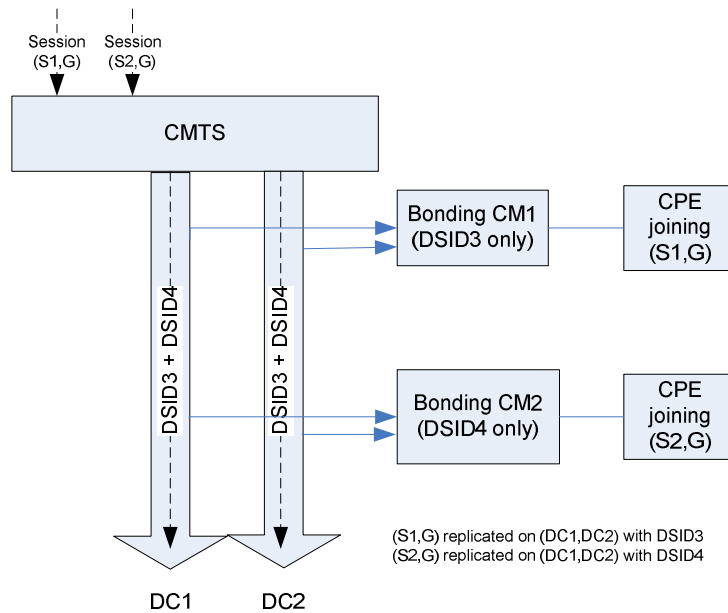


Figure 9-4: DSIDs separate Source-Specific Multicast Sessions

In the example in figure 9-4, the CMTS receives two SSM sessions, (S1,G) and (S2,G) and replicates them both to the downstream bonding group consisting of both DC1 and DC2. By assigning a different DSID to each session, it is able to configure CM1 and CM2 to forward traffic only for the particular SSM session joined by the CPE reached through the CM. The CMTS signals CM1 to recognize DSID3 but not DSID4 and the CMTS signals CM2 to recognize DSID4 but not DSID3. Each CM forwards the proper SSM session traffic and filters the other SSM session traffic based on the DSID.

9.2.2.2 Labeling Multicast Packets with DSIDs

The CMTS MUST label all downstream multicast packets with a DSID. Packets with a known DSID are received by the CM and forwarded to the set of interfaces associated with the DSID. A routing CMTS MUST label traffic for different "IP multicast SSM Channels" or "IP multicast ASM Groups" with different DSIDs, with the exception of well-known IPv6 multicast traffic (refer to clause A.1.2). Thus, with Multicast DSID-based Forwarding, each replication of an (S,G) IP multicast session to a particular DCS is assigned a unique DSID label within a MAC Domain.

A bridging CMTS SHOULD label traffic for different "IP multicast SSM Channels" or "IP multicast ASM Groups" with different DSIDs. If the bridging CMTS is not capable of isolating multicast traffic based on layer-3 information (such as an ASM Group or SSM channel) then the bridging CMTS MUST use different DSIDs for multicast traffic with different destination GMAC addresses.

DSID labeling enables differentiation of multiple replications of an IP Multicast Session (bonded or non-bonded) on downstream channel sets (refer to clause L.2.1.1, Definition of DCS). Hence, it is possible that the CMTS assigns multiple DSIDs to an IP Multicast Session. Non-bonded multicast packets contain a DSID in the DOCSIS header without a sequence number. To prevent a CM from receiving duplicate packets, the CMTS MUST NOT replicate multicast packets labeled with the same DSID on different downstream channel sets that reach the same CM.

The CMTS typically signals only one DSID out of the set of DSIDs that are being used for the replications of a specific IP Multicast Session to the CM. The CMTS has the option to signal multiple DSIDs for the same IP Multicast Session to a CM. However, the CMTS needs to ensure that it does not replicate the IP Multicast session with those DSIDs concurrently on the downstream channel sets reached by the CM. This prevents duplicate delivery of packets to the CM. For example, the CMTS may communicate two DSIDs to the CM, one DSID used for forwarding the stream bonded and another used for forwarding the stream non-bonded, but the CMTS uses only one of those two DSIDs for labeling multicast packets received by a CM. This enables the CMTS to switch a multicast session from bonded to non-bonded without having to incur delay in communicating a DSID.

In order to achieve bandwidth efficiencies, the CMTS SHOULD minimize the number of copies of a multicast packet that need to be delivered on the overall set of downstream channel sets.

9.2.2.2.1 Mixed CM environment

DOCSIS 3.0 networks may contain a mix of DOCSIS 3.0 cable modems and Pre-3.0 DOCSIS cable modems. Pre-3.0 DOCSIS CMs do not support downstream channel bonding. However, the CMTS may need to transmit multicast packets to Multicast Clients behind the Pre-3.0 DOCSIS CMs. A CMTS MUST replicate multicast traffic intended for CMs that do not support Multiple Receive Channels (e.g. DOCSIS 2.0) as non-bonded. It is possible that a given multicast packet is replicated multiple times on a single downstream channel: once non-bonded to be received by CMs that are not receiving multiple downstream channels and one or more times bonded to be received by CMs that are receiving multiple downstream channels.

DSID labeling can ensure that a DOCSIS 3.0 CM does not forward duplicate packets. However, because Pre-3.0 DOCSIS CMs ignore the DSID label on the packet, it is possible that Pre-3.0 DOCSIS CMs receive bonded copies of DSID labeled packets. This may result in Pre-3.0 DOCSIS CMs receiving partial as well as duplicate copies of bonded packets. The CMTS MUST isolate bonded multicast traffic from non-bonded replication on the same downstream channel by transmitting these bonded multicast packets with the Isolation Packet PDU MAC Header (the FC_Type field of the DOCSIS frame set to binary 10). Note that the CMTS MAY transmit bonded multicast traffic with the Packet PDU MAC Header (FC_Type set to 00) when such traffic does not overlap with a non-bonded replication of the multicast session on the same downstream channel. For the replication of non-bonded multicast traffic to CMs with a Frame Control Type Forwarding Capability of 0 (i.e. cannot forward FC_Type 10), the CMTS MUST transmit the non-bonded multicast traffic with the Packet PDU MAC Header (the FC_Type field set to binary 00). Because the CMTS does not know of the CM's capabilities until the CM registers, the CMTS MUST NOT isolate pre-registration IPv6 multicast traffic (clause 9.2.2.2) with the Isolation Packet PDU MAC Header (FC_Type 10).

9.2.2.2.2 Pre-Registration DSID

The Pre-Registration DSID is the DSID for labeling multicast packets used by a CM prior to its completion of the registration process; these multicast packets used for DHCPv6, Neighbor Solicitation (DAD) and IPv6 Router Advertisements, are received only by a CM's IP stack. The CMTS MUST label all link local multicast traffic (as detailed in annex A) with the Pre-Registration DSID.

9.2.2.2.3 Upstream Multicast Traffic from a Multicast Client

According to the requirements in clause 9.1.2.3, when a CM receives upstream multicast packets on its CMCI interface, it forwards packets on its RF Port, the CM IP stack and all of its CMCI ports, except the one from which it received the packets. Additionally, as specified in clause 9.1.2.2, when the CM receives DSID labeled downstream multicast packets, it filters packets from a source MAC addresses which are provisioned or learned as supported CPE devices. Therefore, when forwarding upstream multicast packets to downstream channel sets on the MAC Domain, if the DSID used for those multicast packets is known to the CM from which the packets were received the CMTS MUST NOT alter the source MAC address on those packets. This prevents duplicate delivery of packets to multicast clients behind the CM including the original sender.

9.2.2.3 Communicating DSIDs and group forwarding attributes to a CM

The CMTS is responsible for signalling to the CM a DSID that the multicast traffic is labeled with. The CMTS advertises the Pre-registration DSID value in the MDD message (clause 6.4.28.1.5). The CMTS also communicates DSID values to the CM during the registration process in a REG-RSP message and dynamically using the DBC message after registration.

The CMTS transmits the Pre-registration DSID in the MDD message. The CMTS MUST assign a unique Pre-registration DSID per downstream channel in the MAC domain.

A CMTS is responsible for sending one copy of the IPv6 Well Known (clause A.1.2) and Solicited node traffic to CM and associated CPE devices that require such traffic, which is necessary for DHCPv6, Neighbor Solicitation (DAD) and IPv6 Router Advertisements after registration. The CMTS has the option of continuing to use the Pre-registration DSID for CMs in the operational state or assigning a new DSID for this multicast traffic. In either case, the CMTS MUST include the DSID values for post registration well-known IPv6 traffic for any CM in which Multicast DSID Forwarding has been enabled in the REG-RSP(-MP) message. The CMTS MUST NOT assign a new DSID when it receives a MLD Membership Report for the Solicited Node Address from an IPv6 node that is initializing its stack (traffic associated Neighbor Discovery and Duplicate Address Detection). This allows the CMTS to use the same DSID for all IPv6 provisioning traffic and does not generate a new DSID for each SN address.

The CMTS MUST communicate in the REG-RSP-MP the set of DSIDs for multicast packets to be forwarded by that CM immediately after the registration process.

A CM may have several logical and physical interfaces to internal and external multicast clients. The internal CM IP stack is considered to be a multicast client. Each embedded Service Application Functional Entity (eSAFE) is a potential multicast client connected via a separate logical CPE interface. Each external CPE port is a separate interface to a potential multicast client. For the purpose of IP multicast forwarding, a CM can be thought of as a bridge with one port connecting to the CMTS and up to 16 non-CMTS facing ports connecting to Multicast Clients. These non-CMTS facing ports are henceforth called CMIM-Interfaces because they are enumerated via the CM Interface Mask (CMIM) (refer to clause C.2.1.4.8). The group forwarding attributes associated with a DSID determine a set of interfaces on which the CM forwards downstream multicast packets labeled with that DSID value.

DSID based filtering and forwarding for downstream multicast is triggered by a "JoinMulticastSession" message sent by a Multicast Client, like an IGMP version 2 or 3 or an MLD version 1 or 2 membership report. When the CMTS receives a "JoinMulticastSession" message, it initiates a DBC message to the CM from which the "JoinMulticastSession" was received. The DBC message contains the DSID used for labeling packets belonging to the IP Multicast Session as well as the CMIM and/or Client MAC address(es) of Multicast Client(s) where the multicast packets are to be forwarded. The DBC message optionally contains a SAID if the IP Multicast Session is encrypted. The CM responds to this DBC message after configuring appropriate forwarding rules for the session. After registration of a CM, the CMTS MUST communicate changes to the set of DSIDs used for multicast packets to be forwarded by that CM using DBC messages.

The CMTS tracks the multicast forwarding state established on the CMs via DBC messages and appropriately updates them when Multicast Clients join and leave IP Multicast Sessions. CMs are not aware of the methods used to determine which DSID, downstream channel(s) or Multicast Client MAC addresses are used for transporting a specific IP Multicast Session. These methods are the same whether the CM is in a bonded or non-bonded configuration. The details of processing "JoinMulticastSession" and "LeaveSession" messages depends on the actual protocol used (e.g. IGMPv2 or IGMPv3) and is explained in clause 9.2.5.

9.2.2.4 DSID based Filtering and Forwarding by a Cable Modem

A CM MUST NOT forward downstream multicast packets based on snooped IGMP v2/v3 messages.

Since all multicast traffic that is meant to be forwarded by the CM is labeled with a DSID, the CM MUST discard any multicast packets without a DSID label. The CM discards any packet with an unknown DSID. The CM performs filtering and forwarding of downstream multicast traffic based on DSID values; it does not perform destination GMAC address filtering. The CM MUST NOT discard a multicast packet based on its destination GMAC address. A CM MUST support DSID based multicast forwarding for at least as many DSIDs as reported by the "Multicast Downstream Service ID Support Capability" (see clause C.1.3.1.32).

A mechanism is defined to control multicast packet replication within the CM, as the CM may support multiple egress interfaces. For each DSID, the CMTS specifies the CMIM and/or client MAC addresses of the Multicast Clients intended to receive that IP Multicast Session.

In order to successfully obtain its IP address and register, the CM needs to receive certain multicast packets such as those used for DHCPv6, router discovery and duplicate address detection (see clause A.1.2). Prior to registration, the CM MUST forward to its internal IPv6 and higher stacks all multicast packets received on the RF interface and labeled with the Pre-registration DSID signaled in the MDD message. Prior to registration, the CM MUST discard multicast traffic that is not labeled with the Pre-registration DSID.

The CM only forwards packets labeled with the Pre-registration DSID until it receives a REG-RSP message. The CM MUST discard the Pre-registration DSID prior to adding the DSIDs communicated in the REG-RSP.

The CMTS communicates client MAC addresses based on IGMP/MLD join messages for a particular IP Multicast Session to a CM in a DBC-REQ Message. The CM builds a list of client MAC addresses per DSID using these client MAC addresses. The CM MUST support all learned CPE MAC addresses (clause 9.1.2.1) in its client MAC address list associated with each supported Multicast DSID. In other words, if the number of CPE MAC addresses learned by the CM is 4, then the CM needs to support forwarding of multicast sessions to all 4 CPEs for every Multicast DSID it supports. Thus, if the total number of Multicast DSIDs supported by the CM is 16 then the total number of multicast sessions forwarded by the CM will be $16 \times 4 = 64$.

The CMTS may communicate the CM Interface Mask (CMIM) for static (un-joined) multicast services in which case the Multicast Clients (e.g. embedded STBs) do not explicitly send a "JoinMulticastSession" message. The CM uses the CMIM and client MAC addresses to deduce the set of egress interfaces to which the DSID-labeled multicast traffic is forwarded. If the CMTS signals both the CMIM and Client MAC Address for a DSID then the CM does a logical 'OR' operation.

A CM MUST replicate a DSID labeled multicast packet to only the union of all interfaces identified in the Multicast CM Interface Mask associated with that DSID and all interfaces identified by the list of client MAC addresses associated with that DSID. The upper bound for this union for a DSID is all CM egress interfaces. The CM does not forward multicast packets labeled with a known DSID for which it has no interface defined on which to forward these packets.

A CM MUST replicate a DSID labeled multicast packet only once on each interface. If no Multicast CM Interface Mask or Client MAC Address is configured for the DSID, the CM MUST discard multicast packets labeled with that DSID.

9.2.2.5 Individually Directed Multicast

Individually directed multicast refers to the ability in the DOCSIS network to send a multicast packet on the downstream and ensure that it is forwarded by only one CM rather than the full set of CMs with Multicast Clients that have joined an IP Multicast Session. One potential usage scenario is for IGMPv2/MLDv1 Leave Processing as specified in clause 9.2.5.4.

If the CMTS intends to direct multicast packets to a single CM it should use an individual DSID known only to that CM for such packets.

9.2.3 Downstream Multicast Traffic Encryption

9.2.3.1 Multicast Encryption Overview

When a CMTS encrypts downstream multicast traffic associated with an IP Multicast Session intended to be received and/or forwarded by a group of CMs, it does so with a Security Association (SA) previously signaled to those CMs. This type of Security Association ID (SAID) is defined as Per-Session SAID. A Security Association is said to be "known" at a CM when the CMTS has communicated that SAID in a Security Association Encoding of a MAC Management Message sent to the CM.

A Security Association is not considered to be dedicated to either unicast or broadcast (including multicast) traffic. The CMTS MAY transmit multicast traffic intended for forwarding by a group of CMs with any SA known by those CMs. A Per-Session SAID is unique per a MAC Domain Downstream Service Group (MD-DS-SG).

As described in DOCSIS 3.0 Security Specifications [15], when a CM first authenticates with the CMTS the CMTS provides in its BPI Authorization Response message a Primary SA and (if supported by the CMTS) zero or more Static SAs. A CM's initial BPI authentication may occur immediately after initial ranging in a process called Early Authentication and Encryption (EAE) [15]. If a CM does not perform EAE, it performs its initial BPI authentication immediately after it registers with the CMTS. The Primary SA and Static SAs (if any) established at BPI authentication remain in effect as long as the CM remains authenticated with the CMTS.

DOCSIS versions 1.1 and 2.0 used a mechanism that mapped IPv4 multicast destination addresses to a "dynamic" type Security Association. This mechanism is described in DOCSIS 2.0 BPI Specifications [15] and calls for a CM that recognized an upstream IGMPv2 membership report to send an SA Map Request message to the CMTS. The CMTS responded with an SA Map Reply message that provided an SAID of a "dynamic" type Security Association. The CM then initiated a TEK transaction to obtain the keying material for that dynamic SA.

DOCSIS 3.0 introduces a new mechanism for communicating dynamic SAs for multicast traffic instead of using the SA Map Request and SA Map Reply messages of DOCSIS 1.1/2.0. DOCSIS 3.0 calls for the CMTS to signal to the CM the dynamic Security Association for encrypting downstream multicast traffic in the same MAC Management Message with which it communicates the DSID to the CM for that multicast traffic (see clause 6.4.29).

The CMTS communicates a dynamic Security Association to a CM with a Security Association Encoding (clause C.1.5.5) within a Registration Response (REG-RSP) or Dynamic Bonding Change Request (DBC-REQ) message. Although dynamic Security Associations are primarily intended for encrypting downstream multicast traffic, there is no requirement that they do so. A CMTS MAY encrypt unicast, broadcast or multicast traffic with a Primary, Static or Dynamic SA. A CM is expected to decrypt unicast, broadcast or multicast traffic with the appropriate known SA, regardless of the SA type.

The encryption for multicast sessions can be configured in the Group Encryption Configuration object which is referenced from the Group Configuration Object. The GC entry for a multicast session if configured, points to an entry in the Group Encryption Table. This encryption applies only to joined IP multicast sessions. This includes dynamically joined sessions using multicast management protocol such as IGMP/MLD as well as statically joined sessions using Static Multicast Session Encodings in REG-REQ(-MP) (see clause C.1.1.27). The mechanism by which the CMTS provides encryption for other downstream broadcast and layer 2 multicast traffic is CMTS vendor specific.

Whenever there is a change to the encryption properties configured for a session then the CMTS SHOULD signal the required SAIDs using DBC messages to all the CMs which are listening to that Multicast session.

9.2.3.2 Dynamic Multicast Encryption

The message exchange between the CMTS and the CM for the signalling and initialization of multicast traffic encryption varies depending on the type of multicast session, the capabilities of the modem and the multicast forwarding mode selected by the CMTS. The signalling of Security Associations for encrypted dynamic multicast sessions is described in [15].

9.2.3.3 DSIDs and SAIDs

In general, the set of DSIDs and SAs known at a CM are considered to be independent. The CM is not expected to associate an SA with a DSID. Unless specified otherwise, the CMTS MAY transmit encrypted downstream multicast traffic intended for forwarding by a set of one or more CMs with any combination of an SA known by the CMs and labeled with a DSID known by the CMs. A CM MUST decrypt downstream multicast traffic encrypted with an SA known by the CM and labeled with a DSID known by the CM. A CM silently ignores downstream multicast packets with a known SAID and labeled with an unknown DSID. For example, the CM does not report a key sequence error or CRC error in this case.

When the CMTS replicates a downstream multicast packet onto multiple downstream channel sets of a MAC domain, it labels the replication on each downstream channel set with a different DSID. When the CMTS is configured to map a downstream IP Multicast Session to a specific SA, the CMTS MUST encrypt all replications of the session with that same specified SA.

As detailed in clause G.4.2.2, a CMTS that elects to override a Pre-3.0 DOCSIS CM's DSID Multicast Forwarding mode from GMAC-Explicit(1) to GMAC-Promiscuous(2) has additional requirements for encrypting the multicast traffic that reaches the overridden CM.

9.2.3.4 Pre-Registration Multicast Encryption

Before a CM registers, it receives layer 2 multicasts for DHCP /ARP for IPv4 or DHCPv6/Neighbor Discovery for IPv6. The CMTS labels multicast traffic intended for reception by CMs before registration with a Pre Registration DSID advertised in the MAC Domain Descriptor (MDD) message.

A CM that performs Early Authentication and Encryption (EAE) is provided with at least a Primary SAID and, at the CMTS's option, may also be provided with zero or more Static SAIDs as defined in DOCSIS 3.0 Security Specifications [15]. A CMTS does not encrypt multicast traffic intended to be received by a CM before it completes registration using a Primary or Static SA known at the CM from Early Authentication and Encryption. A CM MUST decrypt downstream multicast traffic received with the Pre-Registration DSID and a known Primary or Static SA prior to its completion of registration.

9.2.4 Static Multicast Session Encodings

The cable operator can configure the cable modem to join IP multicast sessions during registration. Such multicast sessions are called Static Multicast Sessions. The cable operator configures such static multicast sessions using the CMTS Static Multicast Session Encodings (see clause C.1.1.27).

The CMTS MUST communicate in its REG-RSP the DSID used to label packets of the multicast session described by the Static Multicast Group and Source Encoding subtypes in the CMTS Static Multicast Session Encoding. The CMTS MUST include in the DSID Encodings sent in the REG-RSP, a Multicast CMIM subtype with the value of Static Multicast CMIM Encoding it received in the REG-REQ. If the static multicast session is encrypted, the CMTS also communicates in the REG-RSP message the session's SA Descriptor [15].

If the CMTS disables Multicast DSID Forwarding for a CM, the CMTS MUST ignore the CMTS Static Multicast Session Encodings received in the REG-REQ. This implies that the CMTS does not communicate DSIDs and SAIDs to the CM for those CMTS Static Multicast Session Encodings and does not create a multicast replication entry for this CM.

The cable operator can also configure the cable modem to join layer 2 multicast sessions using the Static Multicast MAC Address TLV (see clause C.1.1.23).

9.2.5 IGMP and MLD Support

9.2.5.1 Motivation behind taking CM out of IGMP Control Plane

In DOCSIS 1.1 and 2.0, the cable modem is required to provide IGMP version 2 type snooping functionality in which the CM intercepts IGMP membership reports and establishes forwarding of multicast packets appropriately. Two modes, active and passive are defined. IGMP timers and requirements are specified in [14]. This model has a set of downsides similar to a general purpose Ethernet environment where there is no well defined single point of control.

In the DOCSIS environment the CMTS is a well defined single point of control. Hence, it is desirable that a CMTS control the multicast operations of CMs. This alleviates the need to perform any IPv4 or IPv6 specific multicast operations in the CM and simplifies filtering and forwarding functionality.

Removing the IGMP control plane from the CM offers wide range of benefits as follows:

- Ensures well defined and consistent multicast forwarding behavior in the CM.
- Simplifies the CM since protocol specific knowledge for technologies such as ASM and SSM for IPv4 and IPv6, including the protocols IGMPv3 and MLDv2, is no longer required.
- Easier to incorporate multicast protocol changes since they only affect the CMTS and not the CMs.
- Other multicast protocols like PIM can be supported in the future by utilizing the same CMTS to CM signalling without affecting the CMs.
- It is easier to solve issues related to MAC level aliasing, access and admission control from the CMTS.

9.2.5.2 IP multicast service model support

IGMP for IPv4 and MLD for IPv6 are the two IETF standards based protocols by which CPE devices signal membership for IP multicast Session. While originally intended to be used only by host-type CPEs, they can also be used by router-type CPE devices or CM co-located routers by using IGMP/MLD proxy-routing [57]. IGMP and MLD are the only two CPE multicast membership protocols required to be supported by the CMTS. The CMTS MUST support IGMPv3 [52] and MLDv2 [54].

The membership reports are passed transparently by the CM towards the CMTS. The CMTS operates as an IGMP/MLD querier and as an IPv4/IPv6 multicast router (for a routing CMTS) or snooping switch (for a bridging CMTS). In IPv4 and IPv6 multicast, two service models exist, both of which are supported by DOCSIS 3.0. The "Any Source Multicast" (ASM) model as defined in [29] (for IPv4 but as well applicable to IPv6) and the "Source Specific Multicast" (SSM) model as defined in [58]. In ASM, clients send IGMPv2/v3 or MLDv1/v2 membership reports to "join to an ASM IP multicast group (G)" indicating that they want to receive multicast traffic with any IP source address and the IP multicast destination address G. In SSM, clients send IGMPv3 or MLDv2 membership reports to "subscribe to an SSM IP multicast channel (S,G)" indicating that they want to receive multicast traffic with the IP source address S and the IP multicast destination address G. A CMTS MUST support ASM (Any Source Multicast) as specified in [29] and SSM (Source Specific Multicast) as specified in [58] for both IPv4 and IPv6.

In IGMPv2/MLDv1 [39] [46], each membership report packet contains exactly one JoinMulticastSession for one ASM IP multicast group. Each IGMPv2/MLDv1 membership leave contains exactly one LeaveMulticastSession for one ASM IP multicast group. In IGMPv3/MLDv2 each membership report contains one or more JoinMulticastSession and/or LeaveMulticastSession for ASM IP multicast groups [29] and/or SSM IP multicast channels [58]. Whether or not a particular message element is for an ASM IP multicast group or an SSM IP multicast channel is determined by the multicast group (G) as defined in [29] and [58]. A CMTS MUST forward downstream IPv4 multicast traffic to CPE devices joined through IGMP version 3 [52] "JoinMulticastSession" message element.

NOTE 1: Support for IGMP version 3 includes backward compatibility for IGMP version 2 [39].

A CMTS MUST forward downstream IPv6 multicast traffic to CPE devices joined through MLD version 2 [54].

NOTE 2: Support for MLD version 2 includes backward compatibility for MLD version 1 [46] "JoinMulticastSession" message element.

9.2.5.3 IGMP and MLD membership handling

Multicast Clients send triggered IGMP/MLD membership reports when they want to start or stop receiving an IP Multicast Session. When the CMTS processes these triggered membership reports, the CMTS sends DBC messages to control forwarding of multicast packets by a CM.

When the CMTS receives a JoinMulticastSession message in an IGMP/MLD membership report from the first Multicast Client behind a CM, the CMTS MUST verify if the CM is authorized to receive the IP Multicast Session requested to be joined in the JoinMulticastSession message as described in clause 9.2.7 (IP Multicast Join Authorization). If the CM is authorized, the CMTS MUST send a DBC message to add the DSID along with an SAID (if the session is encrypted) and the Client MAC Address and/or CMIM. If the CM is not authorized, the CMTS MUST NOT send a DBC message to the CM adding the DSID and associated attributes.

When the CMTS receives a subsequent JoinMulticastSession message for the same IP Multicast Session in an IGMP/MLD membership report from a different Multicast Client behind the CM, the CMTS MUST send a DBC message to add the Client MAC Address and/or CMIM for the DSID already communicated to the CM.

Multicast Clients also send periodic IGMP/MLD membership reports when they respond to general queries from the CMTS. These periodic membership reports are important for the CMTS for efficient bandwidth utilization. They are used to overcome the loss of triggered membership reports that would have indicated that a Multicast Client wants to stop receiving an IP Multicast Session. Such a loss may happen if a Multicast Client crashes or reboots or if these membership reports are lost due to problems in the home network. The CMTS MUST track periodic membership reports received from Multicast Clients and time them out as specified in the IGMP/MLD protocol specifications for the IGMP/MLD querier.

When Multicast Clients use IGMPv2/MLDv1 membership reports, they suppress their periodic reports in the presence of simultaneously seen membership reports for the same session from another Multicast Client. This can cause problems with the above mentioned tracking of these membership reports. The CMTS MUST NOT reflect IGMP and MLD membership reports received on the upstream to downstream channel sets.

NOTE: This requirement applies even in the mixed mode environment for DOCSIS 3.0 CMs and Pre-3.0 DOCSIS CMs.

This avoids the report suppression problem and enables tracking of membership reports on a per-CM and per-CMIM-Interface basis. In addition, report suppression helps to provide privacy for membership reports. Reflecting the membership reports to other CMIM-Interfaces and CMs would permit eavesdropping on foreign Multicast Client's join activities.

Membership report suppression does not occur with IGMPv3 and MLDv2. Each Multicast Client interested in an IP Multicast Session will generate membership reports independent of membership reports from other Multicast Clients. Due to this, the CMTS can track IGMPv3/MLDv2 memberships on a per Multicast Client basis. This also simplifies IGMPv3/MLDv2 leave processing.

When the routing CMTS determines that there are no Multicast Clients for an IP Multicast Session behind a CM, the CMTS MUST send a DBC message to delete the DSID associated with that IP Multicast Session. If the bridging CMTS is using a single DSID to forward multiple IP Multicast Sessions, the bridging CMTS MUST send the DBC message to delete the DSID only after all Multicast Clients joined to all IP Multicast Sessions associated with that DSID have either left or not responded to membership reports. When the CMTS determines that a Multicast Client has left an IP Multicast Session, but this is not the last client of this IP Multicast Session behind this CM, the CMTS MUST send a DBC message for the DSID associated with the IP Multicast Session to either remove the Multicast Client's MAC address from the client MAC address list or to update the CMIM, if there is a change in the CMIM.

The CMTS SHOULD NOT forward traffic for an IP Multicast Session on a downstream channel set if no multicast clients are joined to that session on that downstream channel set (subject to any administrative controls).

The CMTS SHOULD NOT send group-specific or group-source-specific IGMPv3/MLDv2 queries in response to IGMPv3/MLDv2 membership reports indicating a leave.

9.2.5.4 IGMPv2/MLDv1 Leave Processing

If there are multiple Multicast Clients on the same egress interface of the CM, periodic IGMPv2/MLDv1 membership reports are subject to suppression. Hence the CMTS needs to send an IGMPv2/MLDv1 group specific query as part of IGMPv2/MLDv1 leave processing ([39] and [46]) to determine if there are any remaining Multicast Clients joined to the same IP Multicast Session. When IGMPv2/MLDv1 leave is received from a Multicast Client behind a CM, it is sufficient to send the IGMPv2/MLDv1 group specific query as an individually directed multicast packet to a specific CM. This minimizes the load on other CMs and is highly desirable from the perspective of maintaining the privacy of IGMPv2/MLDv1 leaves and joins. If the CMTS determines that it needs to send an IGMPv2/MLDv1 group specific query after an IGMPv2/MLDv1 leave is received, the CMTS SHOULD send this query such that it is forwarded only by the CM from which the leave was received by using an individual DSID known only to that CM.

9.2.5.5 IGMP and MLD version and query support

For each CM, the CMTS MUST maintain a highest supported version of IGMP and MLD. The CMTS MUST maintain the IGMP version as v3 for MDF-enabled CMs. The CMTS MUST maintain the MLD version as v2 for MDF-enabled CMs. When the CMTS receives IGMP or MLD membership reports from a CM with a version higher than the maintained version for the CM, then CMTS MUST ignore such reports. As an exception, the CMTS is not required to ignore MLD Membership Reports for Link-Scope Multicast Groups (e.g. Solicited Node Multicast) from a CM with an MLD version of "none" (clause G.4.3.1). For example, if IGMP version for a CM is v2, then IGMPv1 and IGMPv2 membership reports are accepted and IGMPv3 membership reports are silently ignored.

CMTS MAY support mechanisms by which the IGMP or MLD version maintained for a CM can be changed, however these mechanisms are outside the scope of the present document. This mechanism can be used to disable forwarding of multicast traffic through the CM by setting the maintained version to "none" or to work around potential IGMPv3/MLDv2 query compatibility issues in older CPEs by setting the maintained version to "IGMPv2" or "MLDv1".

9.2.5.6 Separation of Query Domains

In a mixed-mode cable environment where CMs in DOCSIS 3.0 mode co-exist with Pre-3.0 DOCSIS CMs, it is important to control which IGMP messages are being forwarded to the CPEs behind the CMs.

It is necessary to prevent forwarding of IGMPv3 membership queries by DOCSIS 1.1/2.0 CMs. DOCSIS 1.1/2.0 CMs are only capable of snooping IGMPv1/v2 messages. If an IGMPv3 membership query would be forwarded to the IGMPv3 capable CPE behind a DOCSIS 2.0/1.1 CM, the CPE would respond with an IGMPv3 membership report. This IGMPv3 membership report would not be recognized by the 1.1/2.0 CM and hence the CM would not be able to properly forward the multicast packets to the CPE. It is also important that the initial join (unsolicited membership report) sent by the CPE also uses IGMPv2. This needs to be controlled by the multicast application and is outside the scope of the present document.

On the other hand, if the Cable Operator wishes to support IGMPv3 and SSM to the CPEs behind 3.0 CMs, the CMTS has to ensure that only IGMPv3 messages are forwarded to the CPE network and IGMPv2 messages are blocked. This is because of the Host Compatibility Mode defined in IGMPv3 [52] which requires a host to switch to the older version of IGMP whenever it receives a query based on the older version.

The CMTS MUST define two separate sets of DSIDs, one for IGMPv2 and another for IGMP v3. These DSIDs are used for the general query messages being sent in the downstream. To enable CMs to receive and forward the IGMP general query messages to all CPE interfaces, the CMTS MUST signal to the CM in the Registration Response a DSID with an appropriate CMIM. In order to prevent forwarding of both IGMPv2 and IGMPv3 General Queries by a single CM, a CMTS MUST NOT signal DSIDs associated with both IGMPv2 and IGMPv3 to a CM at the same time. The CMTS MUST NOT use the same IGMPv2 DSID for IGMPv2 queries being sent on different downstream channel sets. The CMTS MUST NOT use the same IGMPv3 DSID for IGMPv3 queries being sent on different downstream channel sets. Since the IGMPv3 queries are meant to be forwarded by 3.0 CMs only, the CMTS MUST isolate IGMPv3 general query packets from Pre-3.0 DOCSIS CMs by transmitting the IGMPv3 general query packets with the Isolation Packet PDU MAC Header (setting the FC_Type field to 10).

To enable CMs to receive and forward the MLD general query messages to all CPE interfaces, the CMTS MUST signal to the CM in the Registration Response a DSID with an appropriate CMIM. As Pre-3.0 DOCSIS CMs do not support IPv6, there is no DOCSIS 3.0 requirement that the CMTS separate MLDv1 and MLDv2 general queries with a DSID. However, CMTS vendors MAY decide to provide a similar DSID separation of MLDv1 and MLDv2 general queries, as is defined for IGMPv2 and IGMPv3. If the CMTS supports such separation of MLD general queries then the CMTS MUST define two separate sets of DSIDs, one for MLDv1 and another for MLDv2 general query messages. The CMTS MUST NOT use the same MLDv1 DSID for MLDv1 queries being sent on different downstream channel sets within the same MAC domain. The CMTS MUST NOT use the same MLDv2 DSID for MLDv2 queries being sent on different downstream channel sets within the same MAC domain. Since the MLD queries are meant to be forwarded by 3.0 CMs only, the CMTS MUST isolate MLDv1 and MLDv2 general query packets from Pre-3.0 DOCSIS CMs by transmitting the MLD general query packets with the Isolation Packet PDU MAC Header (setting the FC-Type field to 10).

9.2.6 Encrypted Multicast Downstream Forwarding Example

This example involves the forwarding of an encrypted multicast session without PHS to two multicast clients behind a CM. Refer to figure 9-5:

- 1) Multicast traffic labeled with DSID1 is not forwarded through the CM to any of the clients.
- 2) The Multicast Client 1 on Interface A sends out a "JoinMulticastSession" when it wants to join an IP Multicast Session.
- 3) The CM forwards the "JoinMulticastSession" upstream to the CMTS like any other data packet without snooping.
- 4) Assuming the CMTS accepts the joiner, the CMTS selects a DSID and sends a DBC-REQ message that includes the DSID, a client MAC address and a SAID, since the multicast session is encrypted.

NOTE 1: The address in the Client MAC address list is the source address in the "JoinMulticastSession" (i.e. the MAC address of the Multicast Client 1). The CMTS may start sending traffic for that IP Multicast Session labeled with this DSID prior to sending the DBC-REQ message.

- 5) Upon successful reception of a DBC message, the CM adds the DSID to its filter table. In addition, it associates the client MAC address with this DSID in order to correctly forward multicast packets only to the subscribing Multicast Clients. The CM sends DBC-RSP message to the CMTS with appropriate confirmation/error codes.
- 6) CMTS sends a DBC-ACK message after it successfully receives DBC-RSP message from the CM.
- 7) Since the IP Multicast Session is encrypted, the CM sends the TEK-REQ/BPKM Key Request to the CMTS to obtain the TEK key associated for the SAID.
- 8) The CMTS sends TEK key material to the CM in the BPKM Key Reply message.
- 9) When a packet for the IP Multicast Session arrives at the CMTS, the CMTS labels it with the correct DSID, encrypts the packet with the SAID and then forwards it downstream.

- 10) When the multicast packet arrives at the CM, the CM decrypts the packet and only forwards it to interface A on which the Multicast Client 1 is connected (since only Multicast Client 1 is associated with the DSID signaled to the CM).
- 11) The Multicast Client 2 on Interface B of the CM sends out a "JoinMulticastSession" when it wants to join the same IP Multicast Session.
- 12) The CM forwards the "JoinMulticastSession" upstream to the CMTS like any other data packet without snooping.
- 13) Assuming the CMTS accepts the joiner, the CMTS sends a DBC-REQ message that includes the existing DSID for the IP Multicast Session, the second Multicast Client MAC address and the same SAID used for encrypting the Multicast Session.

NOTE 2: The additional Client MAC address is the source MAC address from the "JoinMulticastSession" (i.e. the MAC address of the Multicast Client 2).

- 14) The CM already has the DSID in its filter table. It associates the new client MAC address with this DSID in order to correctly forward multicast packets to the new client.
- 15) The CMTS continues to forward the packets of the IP Multicast Session downstream with the correct DSID label and encrypted with the SAID.
- 16) When the multicast packet arrives at the CM, the CM decrypts the packet and replicates it to interfaces A and B so that both the clients receive the packet.
- 17) The CM sends a DBC-RSP confirming that it received the DBC-REQ.
- 18) The CMTS responds to this message with a DBC-ACK.
- 19) When Multicast Client 1 decides to leave the multicast group, it sends a "LeaveMulticastSession".
- 20) The CM forwards the "LeaveMulticastSession" upstream to the CMTS like any other user data packet without snooping.
- 21) CMTS receives the "LeaveMulticastSession" from Multicast client 1 and sends a DBC-REQ to the CM deleting the MAC address of Multicast Client 1 from the client MAC address list associated with the DSID.
- 22) Upon receiving the DBC-REQ, the CM removes the MAC address of Multicast Client 1 from the client MAC address list associated with the DSID.
- 23) The CMTS continues to forward the packets of the IP Multicast Session downstream with the correct DSID label and encrypted with the SAID.
- 24) When the multicast packet arrives at the CM, the CM only forwards that packet to interface B; so that only Multicast Client 2 receives the packet.
- 25) CM sends a DBC-RSP confirming that it received the DBC-REQ.
- 26) The CMTS responds to this message with a DBC-ACK.
- 27) The second Multicast Client leaves the network without sending a "LeaveMulticastSession".
- 28) After some time, the Membership Timer expires and the CMTS determines that the Multicast Client 2 has left the IP Multicast Session. The CMTS determines this as it did not receive membership reports from Multicast Client 2 during the Membership Timer Interval.
- 29) The CMTS determines that there are no Multicast Clients connected to the CM that are intended to receive the IP Multicast Session. Hence the CMTS sends a DBC-REQ to delete the DSID from the CM.
- 30) When the CM receives the DBC-REQ deleting the DSID, it removes the DSID from its filter table.
- 31) Now when multicast packets arrive at the CM, they will be discarded as the DSID does not match with the set of known DSIDs in the CM.
- 32) CM then sends a DBC-RSP confirming that it received the DBC-REQ.

- 33) The CMTS responds to this message with a DBC-ACK.
- 34) The CMTS continues to forward the packets of the IP Multicast Session downstream with correct DSID label to other CMs.

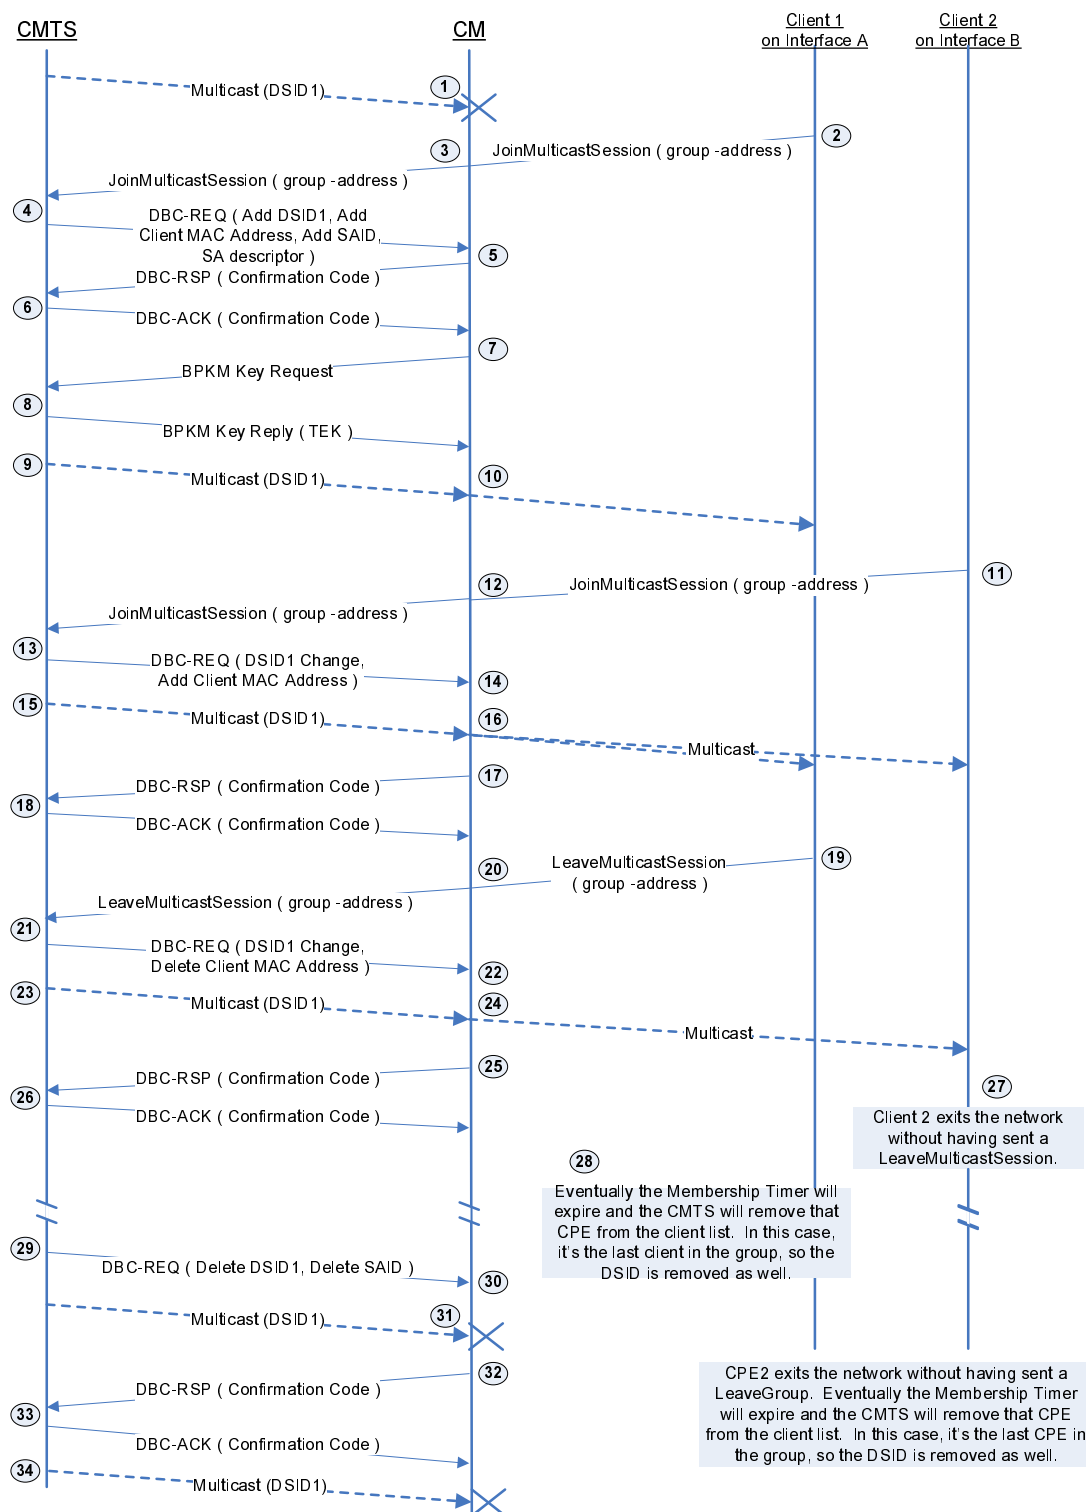


Figure 9-5: Example: Encrypted Downstream Multicast Forwarding

9.2.7 IP Multicast Join Authorization

DOCSIS 3.0 introduces an IP Multicast Join Authorization feature that allows operators to control which IP multicast sessions may be joined by multicast clients reached through a CM. The set of IP multicast clients reached through a CM includes the CM IP host stack itself. This feature controls only the joining of downstream IP multicast sessions; it does not control the ability of any client to transmit IP multicast traffic upstream.

The CMTS enforces IP Multicast join authorization by signalling or not signalling Multicast DSIDs and/or per-session Security Associations. CMTS requirements in this clause for enforcing IP Multicast Join Authorization for CMs that do not implement Multicast DSID Forwarding (e.g. all CM versions before DOCSIS 3.0) and for MDF-disabled CMs require that the operator enable BPI for all such CMs and that the CMTS Group Configuration management table enable per-session encryption. However, it is not necessary to use per-session encryption for enforcing IP Multicast Join Authorization for MDF-enabled CMs because the CMTS controls multicast forwarding by the MDF-enabled CMs by simply signalling or not signalling a DSID used for labeling packets of a multicast session.

The CMTS MUST implement a management object that globally enables or disables IP Multicast Join Authorization Enforcement. When IP Multicast Join Authorization Enforcement is globally enabled, the CMTS MUST NOT enable Multicast DSID Forwarding through a CM for an IP Multicast session that is unauthorized by the IP Multicast Join Authorization feature. When IP Multicast Join Authorization Enforcement is globally disabled, the CMTS MUST authorize all IP multicast joins for all CMs.

The CMTS MUST authorize the following IP multicast sessions to be joined by IP multicast clients reached through a CM:

- IP multicast sessions identified by a Static IP Multicast Session Encoding clause C.1.1.27 in the CM's registration request.
- IPv4 or IPv6 multicast sessions which map to a layer 2 Ethernet multicast MAC address identified in a Static Multicast MAC Address Encoding clause C.1.1.23 in the CM's registration request.
- An IP multicast session for which the highest priority matching "IP Multicast Join Authorization Session Rule" associated with the CM has a "permit" action.
- An IP multicast session that does not match any "IP Multicast Join Authorization Session Rule" associated with a CM when the management object "Default IP Multicast Join Authorization Action" is set to "permit".

The CMTS MUST NOT authorize any IP multicast session not explicitly required to be authorized (as identified in the bulleted list above).

The sessions identified in the first three bullets above are still authorized even if the highest priority matching "IP Multicast Join Authorization Session Rule" associated with the CM has a "deny" action for those sessions.

In order to support the necessary Neighbor Discovery and Duplicate Address detection requirements that IPv6 nodes have, the well known IPv6 addresses and Solicited Node Address traffic are exempt from Multicast Join Authorization enforcement by the CMTS.

The CMTS MUST ignore an IP multicast join request that is not authorized. The CMTS MUST NOT start a new replication or create management objects for an unauthorized join request. The CMTS MUST NOT signal a Multicast DSID to a CM for an IP multicast session that is unauthorized when IP Multicast Join Authorization Enforcement is enabled. The CMTS MUST NOT signal to a CM any security association encrypting an IP multicast session when that session is not authorized for the CM.

DOCSIS 3.0 deprecates the CMTS management object "BPI2 CMTS Multicast Authorization Table", which statically authorizes particular SAIDs to particular CMs. It is replaced with the IP Multicast Join Authorization feature of DOCSIS 3.0. When an operator desires to encrypt IP multicast sessions that are statically joined by CMs, the operator includes a Static IP Multicast Session Encoding in the CM configuration file.

9.2.7.1 Maximum Multicast Sessions

The IP Multicast Join Authorization feature permits an operator to configure the maximum number of multicast sessions joined by clients reached through a CM. Since the CMTS maintains a database of each client for each session on each cable modem, limiting the number of sessions for any one CM can prevent a denial of service attack by a malicious CPE that attempts to exhaust those CMTS resources.

An operator configures a Maximum Multicast Sessions Encoding in the CM configuration file clause C.1.1.27 and the CM includes this encoding in its REG-REQ(-MP) message to the CMTS. This encoding, if specified, limits the maximum number IP multicast sessions that can be dynamically joined (with IGMP or MLD) by clients reached through the CM. The maximum multicast sessions encoding does not restrict the number of statically joined IP multicast sessions. The CMTS **MUST NOT** authorize multicast session join requests that exceed the limit signaled in the Maximum Multicast Sessions Encoding value. A Maximum Multicast Sessions Encoding value of 0 indicates that no dynamic joins are permitted. A Maximum Multicast Sessions Encoding value of 65 535 (the largest valid value) indicates that the CMTS permits any number of sessions to be joined by clients reached through the CM.

If a CM registers with no Maximum Multicast Sessions Encoding, the CMTS **MUST** use the value of a "Default Maximum Multicast Sessions" management object to indicate the maximum number of sessions permitted to be dynamically joined by clients reached through the CM. A Default Maximum Multicast Sessions object value of 65 535 (the largest valid value) configures the CMTS to permit any number of sessions to be joined by clients reached through a CM that does not have an individually configured Maximum Multicast Session Encoding.

9.2.7.2 Session Rules

DOCSIS 3.0 introduces the concept of IP Multicast Join Authorization Session Rules, which are called simply "session rules" in this clause. A session rule applies to a range of IP multicast sessions and identifies whether a multicast client reached through a CM is permitted or denied authorization to join a session within that range.

A session rule can be considered to be a tuple with the members (S prefix, G prefix, priority, action). A session rule applies to a range of IP multicast sessions with sources within the "S prefix" range and destination groups within the "G prefix" range. Both "S prefix" and "G prefix" in a session rule are an IP prefix consisting of an IP address and a "prefix length" with a number of bits from the left. Because more than one session rule can match a particular session, each session rule has a "rule priority" attribute. When a requested IP multicast session for (S,G) matches more than one session rule, the rule with the highest rule priority takes effect. A session rule identifies an authorization "action" that either permits or denies authorization to a particular (S,G) session that matches the rule.

A CMTS **MUST** implement a management object for a "Default IP Multicast Join Authorization Action" with values of "permit" or "deny". When a session join request does not match any session rule, the CMTS **MUST** authorize the join request when the Default IP Multicast Join Authorization Action is "permit". When a session join request does not match any session rule, the CMTS does not authorize the join request when the Default IP Multicast Join Authorization Action is "deny".

In general, an operator selects one of two modes of operation:

- a default to "permit" authorization with session rules that "deny" ranges of session; and
- a default to "deny" authorization with session rules that "permit" ranges of sessions.

A CMTS associates session rules to a CM with two mechanisms:

- IP Multicast Profiles; and
- static IP Multicast Session Rules.

The IP Multicast Join Authorization Encoding in a CM configuration file specifies both mechanisms to the CMTS. The CMTS searches all session rules associated with a CM to find the highest priority rule matching an IP multicast join request.

9.2.7.2.1 IP Multicast Profiles

At the CMTS, an operator configures a named "IP Multicast Profile" with a set of IP Multicast Join Authorization Session Rules.

The IP Multicast Join Authorization Encoding in the CM configuration file clause C.1.1.18.1.9 provides the name of one or more IP Multicast Profiles. The CMTS associates with the CM the union of all session rule sets configured for the IP Multicast Profiles named in this encoding. The CMTS **MUST** support at least 2 Join Authorization Rules per IP Multicast profile and **SHOULD** support at least 16 Join Authorization Rules per IP Multicast profile.

9.2.7.2.2 Static IP Multicast Join Authorization Rules

The IP Multicast Join Authorization Encoding also can contain explicit, static IP Multicast Join Authorization Rules. The CMTS associates with the CM all static session rules defined in the encoding.

9.2.7.3 CM Configuration File

An IP Multicast Join Authorization Encoding clause C.1.1.18.1.9 in the CM configuration file and CM registration request determines the set of IP Multicast Join Authorization Session Rules associated with the CM. Because the IP Multicast Join Authorization encoding is a subtype of the TLV-43 DOCSIS Extension Information encoding, CMs operating at all DOCSIS versions will include the encoding in a registration request message to the CMTS.

The IP Multicast Join Authorization Encoding includes subtypes that define:

- an "IP Multicast Profile Name" that identifies a list of multicast session rules configured in the CMTS;
- "static IP Multicast Session Rules", each of which directly defines a single IP multicast session rule; and/or
- the "Maximum Multicast Sessions" permitted to be dynamically joined by clients reached through the CM.

9.2.7.3.1 IP Multicast Profile Name Subtype

A CMTS **MUST** accept an IP Multicast Profile Name subtype in an IP Multicast Join Authorization Encoding as identifying a set of session rules configured at the CMTS to be associated with the CM. The CMTS **MUST** accept at least 16 IP Multicast Profile Name encodings for a single CM. The total number of IP Multicast Profiles supported in a CMTS is vendor specific. If a CM registers with more IP Multicast Profile Names than are supported by the CMTS, the CMTS **MUST** ignore the additional profiles, as ordered in the REG-REQ(-MP). If the REG-REQ (-MP) message does not contain a Multicast Profile Name sub-encoding, then the CMTS **MUST** associate with the CM a configured Default Multicast Profile Name List.

In order to avoid requiring an operator to simultaneously update the configuration of all CMTSs and CMs in a region, a CMTS **MUST** support registration of CMs that reference an IP Multicast Profile Name that is not yet configured in the CMTS. When a CM registers with an unconfigured IP Multicast Profile Name, the CMTS **MUST** automatically create an IP Multicast Profile with that profile name and containing no session rules. When the CMTS automatically creates an IP Multicast Profile, the CMTS **MUST** signal an "informational" severity log message.

9.2.7.3.2 Static IP Multicast Session Rule Subtype

A CMTS **MAY** accept Static IP Multicast Session Rule subtypes in an IP Multicast Join Authorization Encoding as defining session rules associated with the CM. If a CMTS does not accept Static IP Multicast Session Rule subtypes, the CMTS **MUST** silently ignore the encoding. If supported, the CMTS **MUST** support at least 16 IP Multicast Join Authorization Static Session Rules for each CM.

If supported, the CMTS maintains a management object that reports for each CM an IP Multicast Static Session Rule List learned from that CM in its REG-REQ(-MP). The CMTS **MAY** recognize when multiple CMs report the same contents of IP Multicast Join Authorization Static Session Rules and so can refer to the same Static Session Rule List ID. The CMTS assigns an IP Multicast Join Authorization Static Session Rule List Identifier to each unique set of IP Multicast Join Authorization Static Session Rules. The minimum number of different IP Multicast Join Authorization Static Session Rule lists supported by a CMTS is vendor-specific.

9.2.7.4 Matching Session Rules

The CMTS **MUST** associate with a CM all session rules in the IP Multicast Profile Name encodings referenced in the CM's registration request. In addition, if the CMTS accepts Static IP Multicast Join Authorization Session Rule subtypes, the CMTS **MUST** also associate with the CM the static session rules signaled in the CM's registration request. The CMTS matches the requested IP multicast session with one or more session rules when the source S is within the source prefix and the group G is within the group prefix of the session rule. When more than one session rule is matched, the CMTS selects the matching session rule with the highest rule priority. The CMTS uses the "action" of the highest priority matching session rule to determine whether the session is authorized. If no session rule matches the join request, the CMTS uses the configured Default IP Multicast Join Authorization Action. If more than one matching session rule has the same highest priority, the particular session rule selected by the CMTS is vendor-specific.

A CMTS receives join requests that are for either source-specific-multicast (SSM) sessions or for any-source-multicast (ASM) sessions. The CMTS MUST match a join request for an SSM session (S,G) to a session rule when both the source S matches the S prefix and the destination group G matches the G prefix of the session rule.

A CMTS MUST match a join request for an Any-Source-Multicast (ASM) group to (G) to a session rule that contains a G prefix field that includes the requested group G and an S prefix field of the session rule matches all sources (i.e. a source prefix length of zero bits). A CMTS MAY map ASM membership reports received from IP multicast clients to SSM sessions received on a network system interface. If the CMTS maps an ASM join request to (G) to an SSM session (S,G), the CMTS MUST match only session rules for which the mapped-to SSM session source S is within the S prefix field of the session rule.

9.2.7.5 IP Multicast Profile Changes

A CMTS MUST support changes to the set of session rules associated with an IP Multicast Profile while a CM remains registered that references that IP Multicast Profile Name. A CMTS MUST apply an updated IP Multicast Profile to subsequent join requests from clients reached through a CM that references the profile. For example, when a CMTS is configured to add new session rules to an IP Multicast Profile, the CMTS includes those rules for subsequent join requests from an already-registered CM that referenced the IP Multicast Profile Name.

When the CMTS configuration of session rules for an IP Multicast Profile changes such that all IP multicast sessions forwarded through a CM using a Multicast DSID are no longer authorized, the CMTS SHOULD dynamically delete on the CM that Multicast DSID and/or security association for the session. A CMTS deletes a security association on an MDF-enabled CM by sending a DBC-REQ to delete the security association. A CMTS deletes a security association on an MDF-disabled or MDF-incapable CM by sending a TEK Invalid BPI key management message [15].

When the CMTS configuration of session rules for an IP Multicast Profile changes such that no CMs reached by a particular replication of an IP multicast session on a downstream channel set remain authorized, the CMTS SHOULD discontinue the replication of the IP multicast session on that downstream channel set.

The CMTS MUST support deletion of an IP Multicast Profile while a CM remains registered that references the profile. The CMTS MUST NOT match session rules for a deleted profile for IP multicast sessions subsequently joined by CMs referencing the deleted profile. When the deletion of an IP Multicast Profile results such that all IP multicast sessions forwarded through a CM using a Multicast DSID are no longer authorized, the CMTS SHOULD dynamically delete on the CM that Multicast DSID and/or security association for the session.

When the deletion of an IP Multicast Profile results such that no CMs reached by a particular replication of an IP multicast session on a downstream channel set remain authorized, the CMTS SHOULD discontinue the replication of the IP multicast session on that downstream channel set.

10 Cable Modem - CMTS Interaction

10.1 CMTS Initialization

The mechanism utilized for CMTS initialization (local terminal, file download, SNMP, etc.) is described in [9]. The CMTS meets 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, the CMTS MUST NOT generate any downstream messages (including SYNCs and UCDs). This will prevent CMs from transmitting.
- The CMTS MUST provide the information defined in clause 6 to CMs for each upstream channel.

10.2 Cable Modem Initialization and Reinitialization

A cable modem **MUST** initialize or reinitialize as shown in figure 10-1. This figure shows the overall flow between the stages of initialization in a CM. The more detailed finite state machine representations of the individual stages (including error paths) are shown in the subsequent figures. Timeout values are defined in annex B.

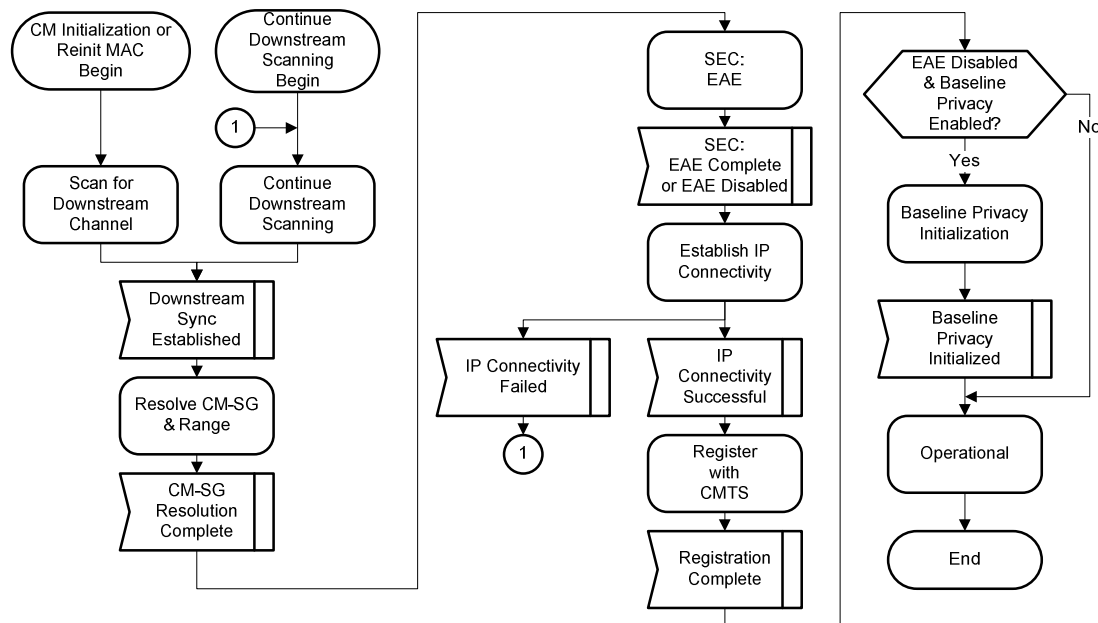


Figure 10-1: Cable Modem Initialization Overview

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 (including scanning continuation when necessary).
- Service group determination and ranging.
- Authentication.
- Establish IP connectivity.
- Registration.

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 [15] (e.g. X.509 certificate) used to authenticate the CM to the security server and authenticate the responses from the security and provisioning servers.

10.2.1 Scan for Downstream Channel

On initial power-on the CM **MUST** set its CM initialization reason to **POWER_ON**. On initialization or a "Reinitialize MAC" operation, the cable modem **MUST** acquire a Primary-Capable downstream channel. The CM **MUST** have non-volatile storage in which the last operational parameters are stored. Unless directed otherwise, the CM **MUST** first try to re-acquire the downstream channel from non-volatile storage. If this fails, the CM **MUST** begin to continuously scan the channels of the entire downstream frequency band of operation until it finds a valid primary downstream signal.

A downstream signal is considered to be valid for use as a CM's Primary Downstream Channel when the modem has achieved the following steps:

- synchronization of the Physical Media Dependent and Transmission Convergence sublayers as defined in [12];
- recognition of SYNC downstream MAC messages.

While scanning, it is desirable to give an indication to the user that the CM is doing so (see [10]).

The Downstream Active Channel List TLV (if provided) from an MDD message received on a non-primary-capable downstream channel may be used by the CM as a "hint" in locating a primary-capable downstream channel.

Once a candidate Primary Downstream Channel has been identified, the CM SHOULD remember where it left off in the scanning process so that it may continue where it left off if necessary.

10.2.2 Continue Downstream Scanning

When the CM determines that the current candidate primary channel is unsuitable, the CM MUST resume scanning downstream spectrum for a suitable candidate downstream channel. The CM SHOULD continue to scan the previously unscanned spectrum.

10.2.3 Service Group Discovery and Initial Ranging

The CMTS needs to determine the service group of a DOCSIS 3.0 CM for channel bonding and load balancing purposes. As described in figure 10-2, the CM MUST attempt to determine its MAC Domain Downstream Service Group ID (MD-DS-SG-ID) if an MDD is present on the downstream. If successful, the CM MUST provide the MD-DS-SG-ID it has selected to the CMTS in the Bonded Initial Ranging Request (B-INIT-RNG-REQ) message. If the CM could not determine its MD-DS-SG-ID then it MUST send a B-INIT-RNG-REQ with the MD-DS-SG-ID set to zero. The CMTS replies to the B-INIT-RNG-REQ with a RNG-RSP message. In order to resolve the upstream service group (MD-US-SG) associated with this CM, the CMTS may include an Upstream Channel Adjustment in this RNG-RSP message. If this occurs, the CM MUST tune to the new channel and sends an Initial Ranging Request (INIT-RNG-REQ) message. The CMTS responds with a RNG-RSP message, possibly including another Upstream Channel Adjustment.

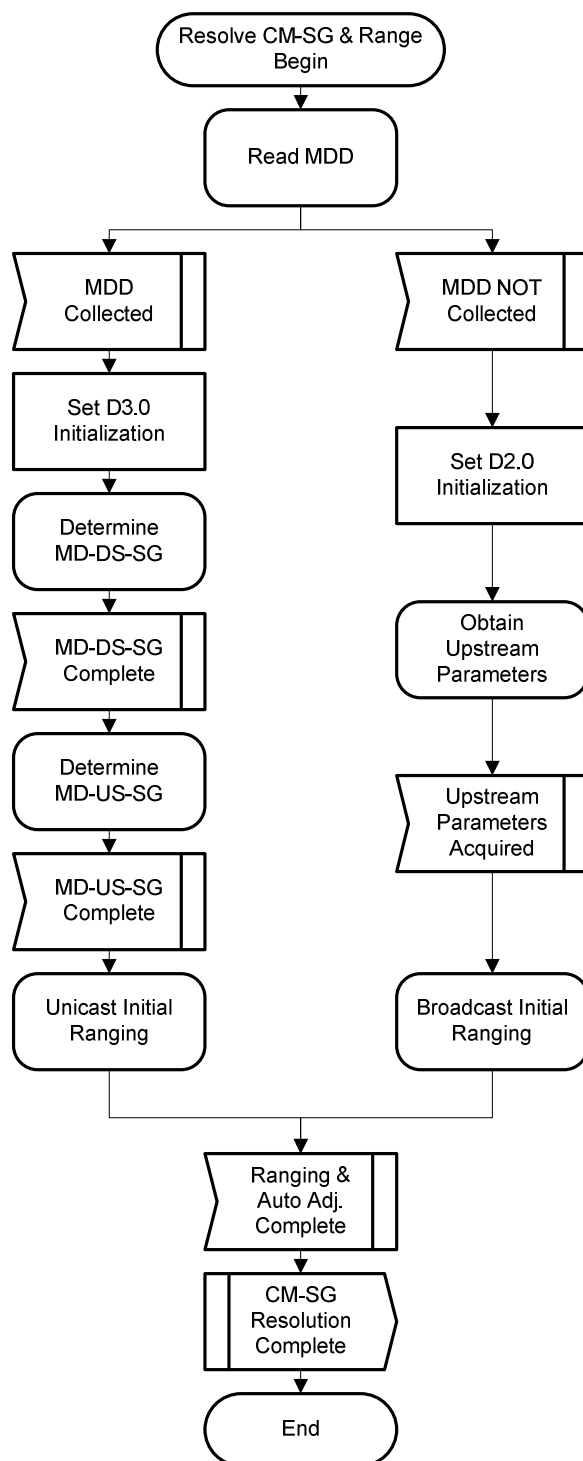


Figure 10-2: Resolve Service Group (SG) and Range

If the determination of the upstream or downstream service group is unsuccessful then the CM continues to gather upstream parameters and then range before continuing on to establish IP connectivity and then register with the CMTS. This process follows very closely what a CM would do on a DOCSIS 2.0 or earlier CMTS.

10.2.3.1 Read MAC Domain Descriptor (MDD)

A DOCSIS 3.0 Compliant CMTS periodically transmits a MAC Domain Descriptor (MDD) MAC management message on all Downstream Channels of a DOCSIS 3.0 MAC Domain. On non-Primary-Capable Channels, the CMTS transmits a MDD message that contains at least the MDD Header with the Downstream Channel ID on which the MDD is sent. On Primary-Capable Channels, the CMTS transmits a MDD message which contains the MDD header as well as TLVs and sub-TLVs containing the following information:

- Information for each Downstream Service Group comprised of the MD-DS-SG-ID along with the set of DCIDs that comprise the MD-DS-SG.
- Channel parameters (e.g. frequency, modulation) for each downstream channel in the MAC Domain as well as an indication of whether that channel is Primary Capable.
- Upstream Active Channel List.
- Upstream Ambiguity Resolution Channel List.
- Upstream Frequency Range.
- Downstream Ambiguity Resolution Frequency List containing a list of downstream frequencies that the CM uses to resolve the MD-DS-SG-ID.
- Other information which is not relevant for the service group determination but which is utilized in later stages of the initialization process.

In certain circumstances, the CM could receive multiple MDD messages with different source MAC addresses, in which case the CM **MUST** attempt to use the MDD message, which is valid for a primary-capable downstream channel. The CM can validate the MDD message TLVs or it can collect MDD messages with the source MAC address learned from the SYNC message.

If, for any reason, the MDD message becomes too long to fit within a single MAC management message, the CMTS fragments the MDD message as described in clause 6.4.28.

The CM **MUST** attempt to read the MDD message from the downstream channel as shown in figure 10-3.

- 1) The CM starts its Lost MDD timeout Timer.
- 2) The CM waits for the arrival of MDD message fragments.
- 3) If the MAC address of the CMTS MAC Domain is not already known, then the CM stores the source MAC address of the received MDD fragment as the MAC address of the MAC domain and adds the fragment to the collection of fragments. At this point the MAC address of the CMTS MAC domain is considered to be known.
- 4) If the MAC address of the CMTS MAC Domain is already known, then upon receiving an MDD message fragment, the CM compares the source MAC address of the newly collected MDD fragment against the known MAC address of the MAC Domain. If the MAC addresses do not match, then the CM discards the fragment and awaits another fragment. If the MAC addresses match, then the fragment is added to those already collected.
- 5) Any time that the CM collects another MDD fragment, the CM **MUST** first check to see whether the change count has been incremented. If the change count has been incremented, then the CM **MUST** discard all collected fragments with the old change count. In either case, the CM then checks to determine whether the entire MDD message has been collected. If it has then the CM ends this process. If all of the fragments of the MDD message have not been collected, then the CM returns to step 2.
- 6) If the Lost MDD timeout Timer expires before the entire MDD message has been collected then the CM informs the calling process of the failure to collect an MDD and exits this process.

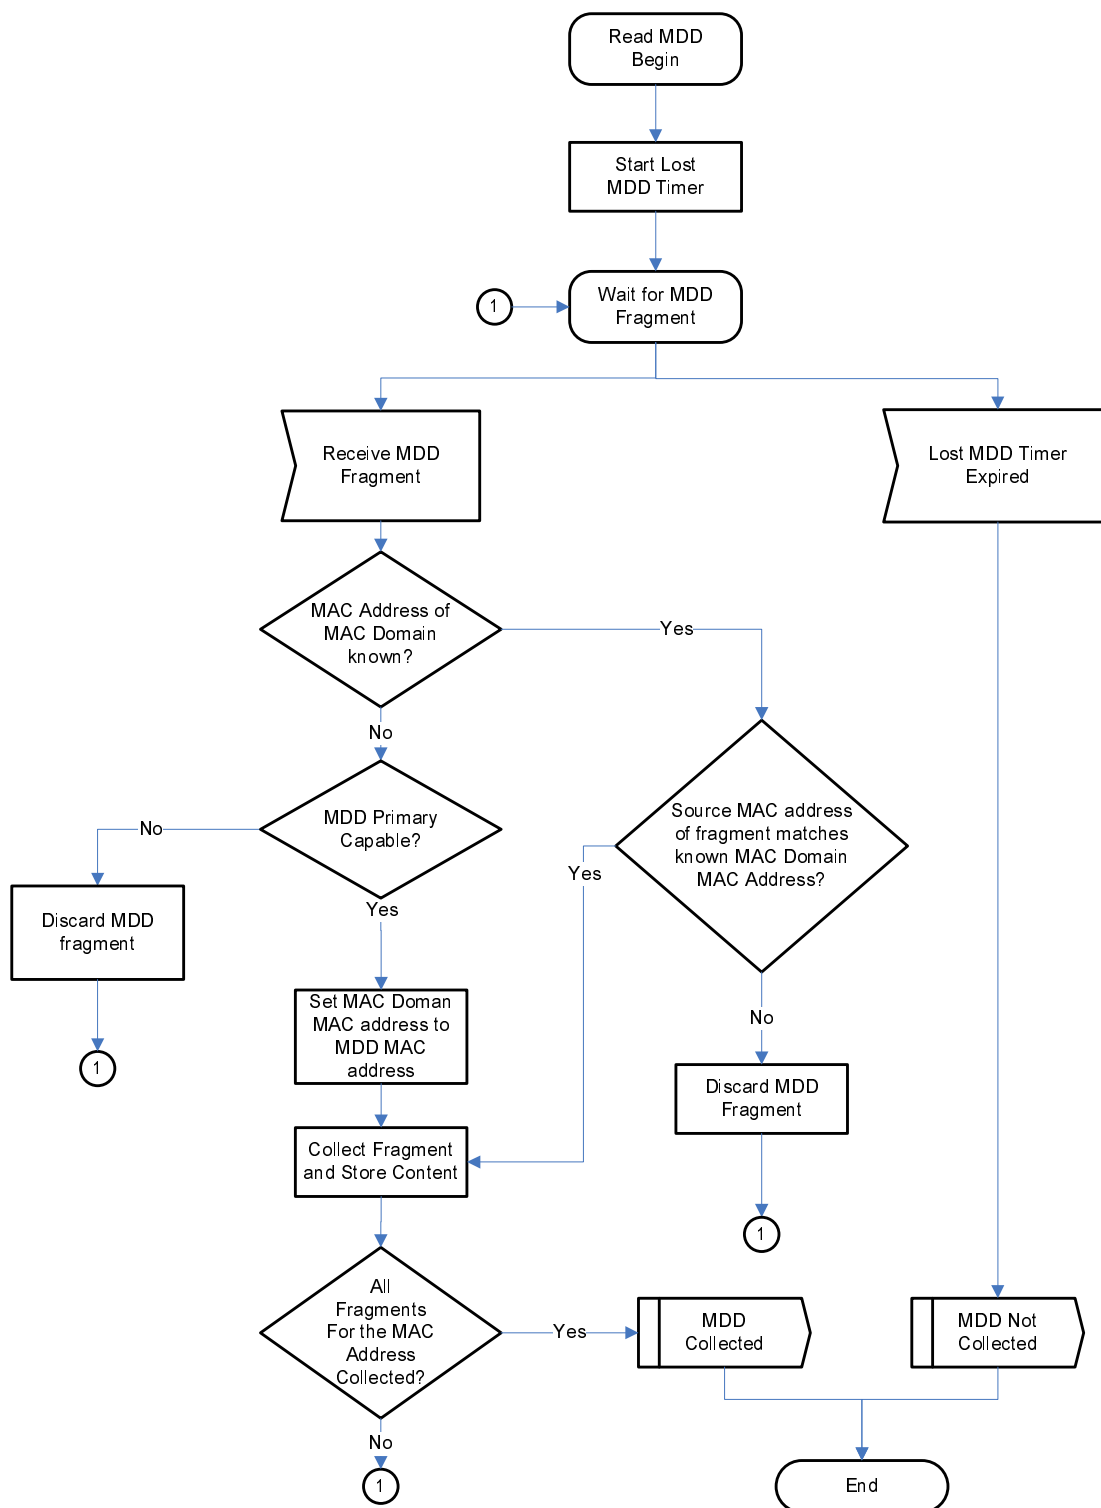


Figure 10-3: Read MAC Domain Descriptor (MDD)

10.2.3.2 MDDs Not Found on Primary Downstream

If no MDDs are detected on the candidate Primary Downstream Channel, the CM MUST revert to DOCSIS 2.0 operation and continue to gather upstream parameters and then range before continuing on to establish IP connectivity and then register with the CMTS (see figure 10-2).

10.2.3.3 Determination of MD-DS-SG

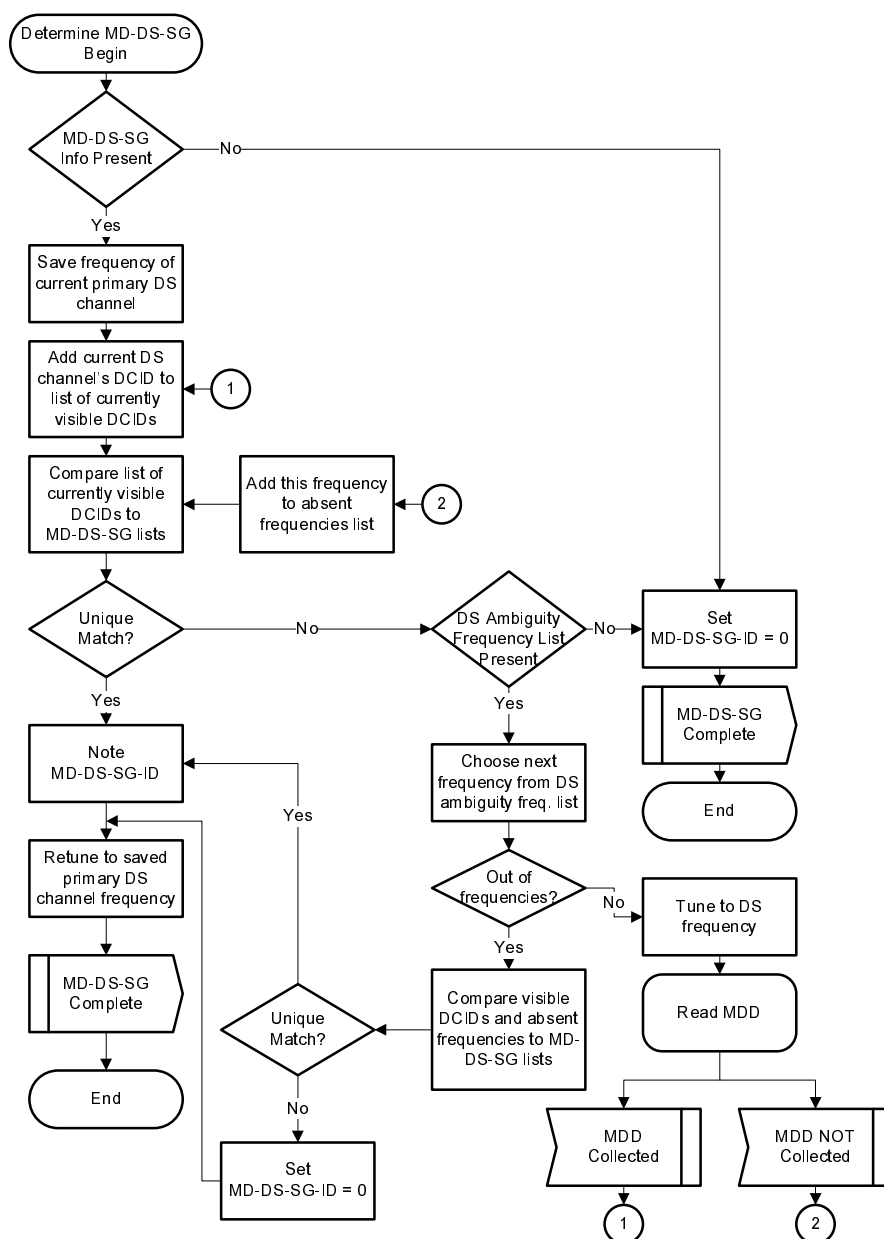


Figure 10-4: Determine MD-DS-SG

The CM MUST attempt to determine its MD-DS-SG according to figure 10-4. This process is described as follows:

NOTE: The CM keeps track of the "list of currently visible DCIDs" by accumulating a list of all DCID values within the MAC Domain of the primary Downstream Channel that it encounters while following the process described in figure 10-4. This "list of currently visible DCIDs" is used to determine the proper MD-DS-SG for the CM.

- 1) If the Primary Downstream Channel's MDD message did not contain at least one MAC Domain Downstream Service Group (MD-DS-SG) TLV, then the CM sets its MD-DS-SG ID to zero and exits downstream service group resolution.
- 2) The CM stores the frequency of the current (primary) DS channel. Then the CM reads the current DCID from the MDD message and adds the DCID to the list of currently visible DCIDs.
- 3) The CM constructs a list of "candidate" MD-DS-SGs. A "candidate" MD-DS-SG is one which is listed in the Primary Downstream Channel's MDD and which contains the current channel's DCID.

- 4) If the list of "candidate" MD-DS-SGs contains a single element, then the MD-DS-SG ID is noted and MD-DS-SG resolution is complete.
- 5) If there is not a unique match but the CM finds that the Downstream Ambiguity Resolution Frequency List TLV is not present in the Primary Downstream Channel's MDD message, then the CM sets its MD-DS-SG ID to zero and exits downstream service group resolution.
- 6) If the Downstream Ambiguity Resolution Frequency List TLV is present in the MDD message and if the list of candidate MD-DS-SGs contains more than one MD-DS-SG, then the CM tunes to the next frequency listed in the Downstream Ambiguity Resolution Frequency List TLV and attempts to read an MDD (see clause 10.2.3.1, read MDD). If an MDD message is found on the new channel, the CM reads the new channel's DCID from the MDD message and adds the new DCID to the "list of currently visible DCIDs". The CM then constructs a new list of "candidate" MD-DS-SGs. In this case, a "candidate" MD-DS-SG is one which is listed in the original channel's MDD and which contains all of the DCIDs from the "list of currently visible DCIDs". If this "candidate" list contains a single entry then the MD-DS-SG-ID is noted, the CM retunes the receiver to the original primary downstream frequency and MD-DS-SG resolution is complete.
- 7) If the MDD is not successfully obtained on the new channel, the CM adds the frequency to its "absent frequencies list". If the Downstream Ambiguity Resolution Frequency List contains more frequencies, then the CM repeats step 6; otherwise, it continues to step 8.
- 8) If the CM runs out of frequencies in the Downstream Ambiguity Resolution Frequency List TLV, but the candidate list of MD-DS-SGs still contains more than one element, the CM attempts to narrow the list further by incorporating the "absent frequencies list". Any candidate MD-DS-SG containing a channel at a frequency included in the "absent frequencies list" is eliminated from the candidate MD-DS-SG list. After this step, if the candidate list of MD-DS-SGs contains exactly one element, the MD-DS-SG ID is noted, the CM retunes the receiver to the original primary downstream frequency and MD-DS-SG resolution is complete. If the CM fails to retune the receiver the original primary downstream frequency then the CM MUST continue scanning the downstream spectrum for a new candidate primary downstream channel.
- 9) If the candidate list of MD-DS-SGs does not contain exactly one element after step 8 has been completed, then the CM exits downstream service group resolution and sets its MD-DS-SG ID to zero.

10.2.3.4 Ranging Holdoff

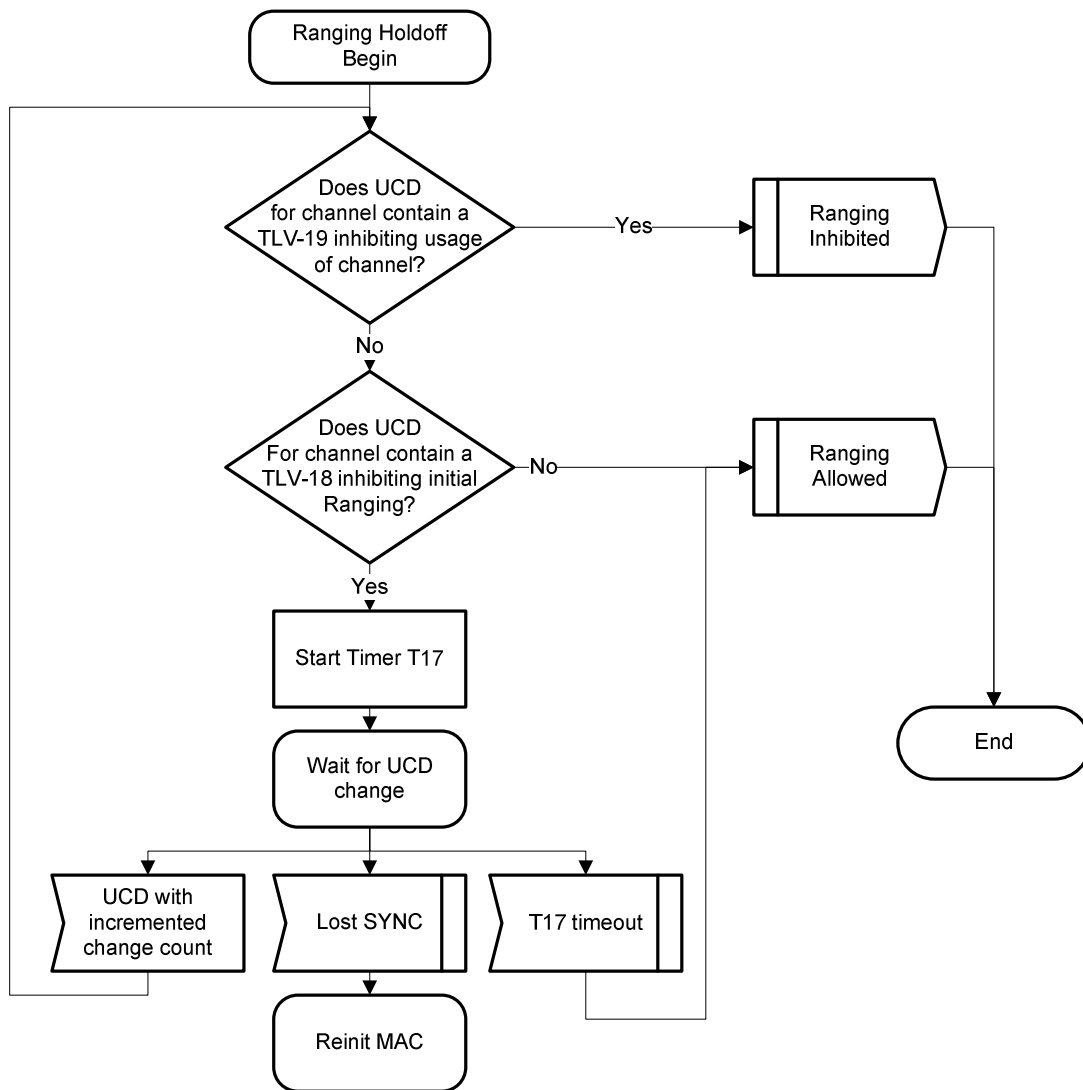


Figure 10-5: Ranging Holdoff

The CM **MUST** check for ranging holdoff direction per figure 10-5 prior to sending an initial ranging request message when it performs initial maintenance for any of the following reasons:

- Power on initialization.
- Reinitialize MAC event, except when triggered by a DCC-REQ with an initialization technique of zero.
- Upstream Channel ID Configuration Setting in configuration file.
- Downstream Frequency Configuration Setting in configuration file.
- Upstream Channel Id Override in RNG-RSP.
- Downstream Frequency Override in RNG-RSP.
- UCD change prior to having sent at least one initial ranging request message (can restart the T17 timer as described below).

The CM MUST NOT check for ranging holdoff direction when it performs initial maintenance for any other reason. Some examples of this include:

- DBC-REQ.
- UCD change with ranging required TLV after having sent at least one initial ranging request message.
- DCC-REQ (applies to all initialization techniques).
- REG-RSP (with TCC).
- RNG-RSP with Upstream Channel Adjustment TLV.

The following rules describe the ranging holdoff operation:

- 1) After selecting an upstream channel for initial ranging, the CM MUST extract the parameters for this upstream from the UCD. If the UCD message contains a Type 19 TLV, the CM MUST (except as described above) perform a bitwise AND of its Ranging Class ID with the TLV 19 Value. If the result of the bitwise AND is zero, the CM MUST consider the channel unusable and try other channels until it finds a usable channel.
- 2) If the UCD contains a Type 18 TLV, the CM MUST (except as described above) perform a bitwise AND of its Ranging Class ID with the TLV-18 Value. If the result of the bitwise AND is equal to the CM's Ranging Class ID, the CM MUST inhibit initial ranging and start the T17 timer. If the UCD Change Count in the UCD message for the channel is incremented while the T17 timer is active, the CM will re-inspect the TLV-18 and TLV-19 value and re-start the T17 timer if necessary. If the T17 timer expires or the TLV-18 value is updated to permit ranging for the CM's Ranging Class, the CM will resume the ranging process. If the CM should undergo a Lost SYNC event while waiting for T17, it MUST reinitialize the MAC with a CM Initialization Reason of T17_LOST_SYNC.
- 3) After having transmitted at least one Initial Maintenance RNG-REQ message, the CM MUST ignore TLV-18 or TLV-19 values in any new UCD message for the channel even if the new UCD contains a Ranging Required TLV.

10.2.3.5 Determination of MD-US-SG

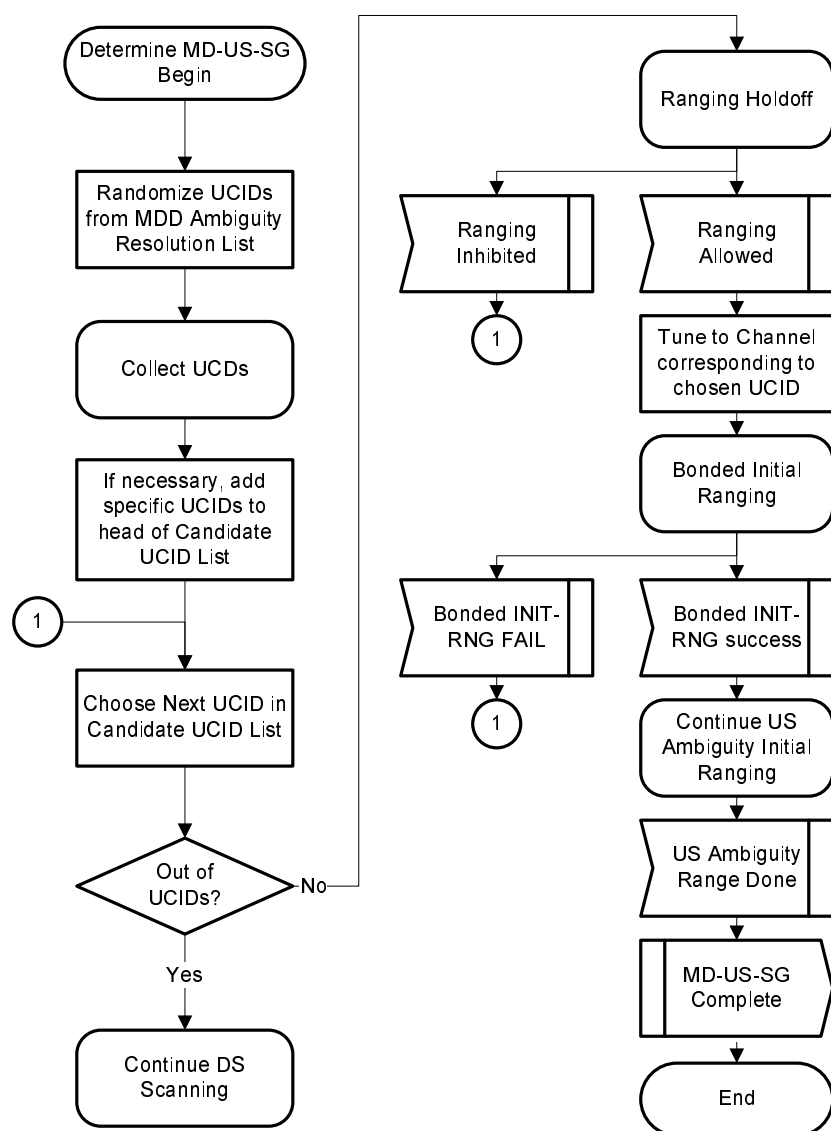


Figure 10-6: Determine MD-US-SG

The following clauses and figure 10-6 explain the steps that a CM MUST perform in order to resolve MD-US-SG resolution.

The CM MUST store the UCID and the transmit power level of all the US channels in its latest operational Transmit Channel Set in non-volatile memory.

Refer to figure 10-6:

- 1) Based on the MDD message received on its Primary Downstream Channel, the CM creates a "Candidate UCID list" by randomly ordering the list of UCIDs in the Upstream Ambiguity Resolution Channel List. If any of these upstream channel UCIDs are stored in non-volatile memory as the last operational transmit channel set, then the CM SHOULD move these UCIDs to the front of the list while maintaining their random ordering relative to each other. In addition, if a specific UCID was sent in an Upstream Channel ID Override TLV in a RNG-RSP message, an Upstream Channel ID Configuration TLV in the CM Configuration File or an Upstream Channel ID TLV in a DCC-REQ message, the CM adds this UCID to the head of the "Candidate UCID List".

- 2) The CM now reads UCD messages and finds the PHY parameters for the upstream channels with UCIDs listed in the "candidate UCID List". If timer T1 expires and the CM has not received any valid UCD messages it **MUST** continue scanning.
- 3) Taking the UCID at the head of the candidate UCID list, the CM performs Ranging Holdoff processing and configures the transmitter for that channel and attempts Bonded Initial Ranging. If the channel was stored in non-volatile memory then the CM **SHOULD** transmit using the stored transmit power level for that channel. If this ranging process fails, then the CM repeats the process with the next UCID in the ordered list.
- 4) Once Bonded Initial Ranging succeeds, the CM continues upstream ambiguity resolution initial ranging as directed by the CMTS.

10.2.3.5.1 Bonded Initial Ranging

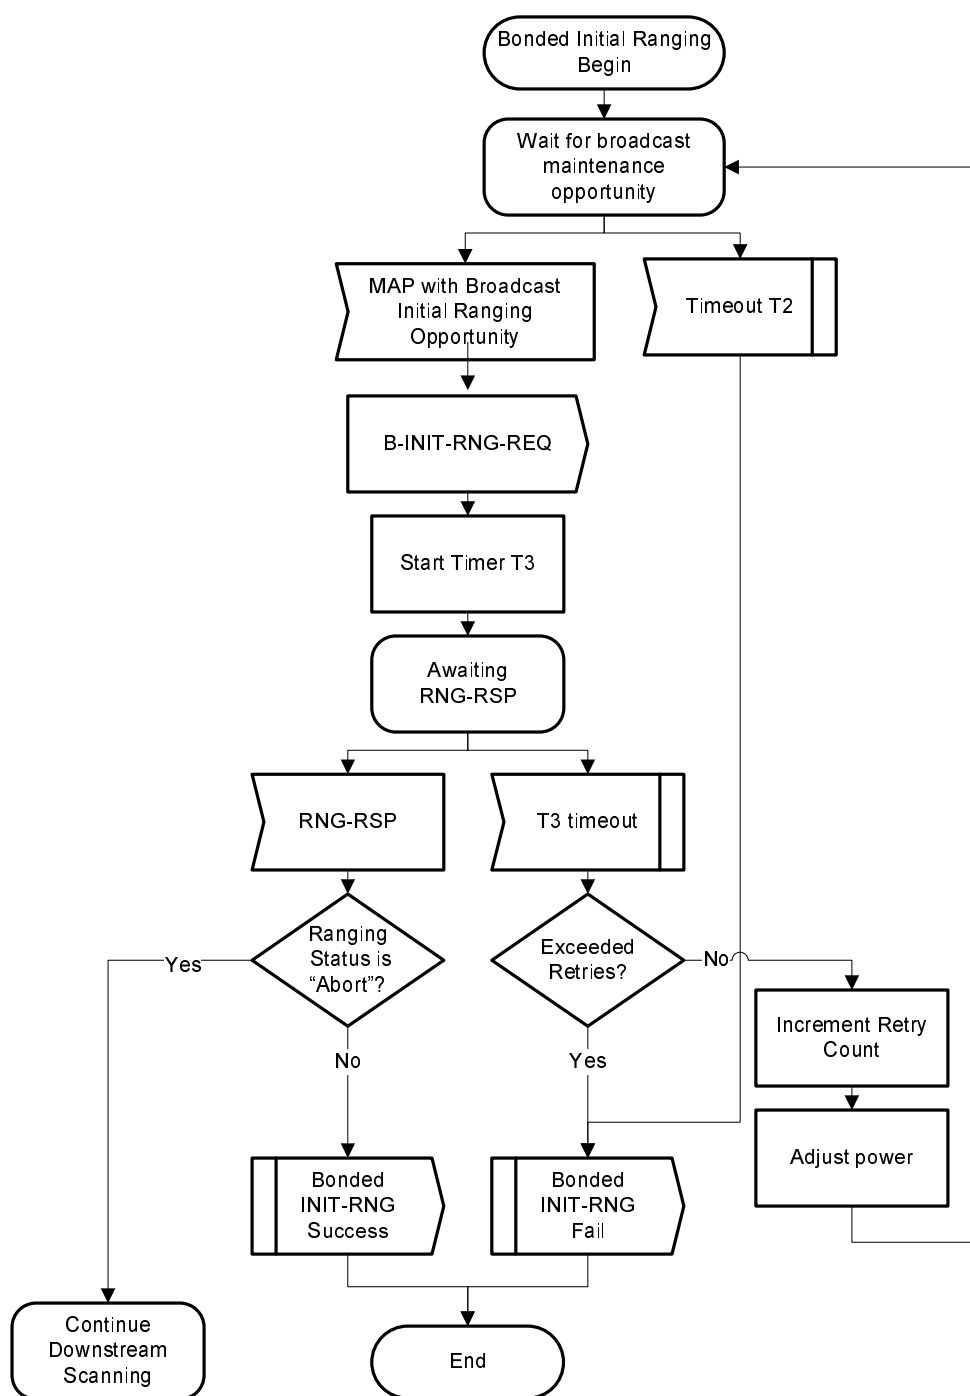


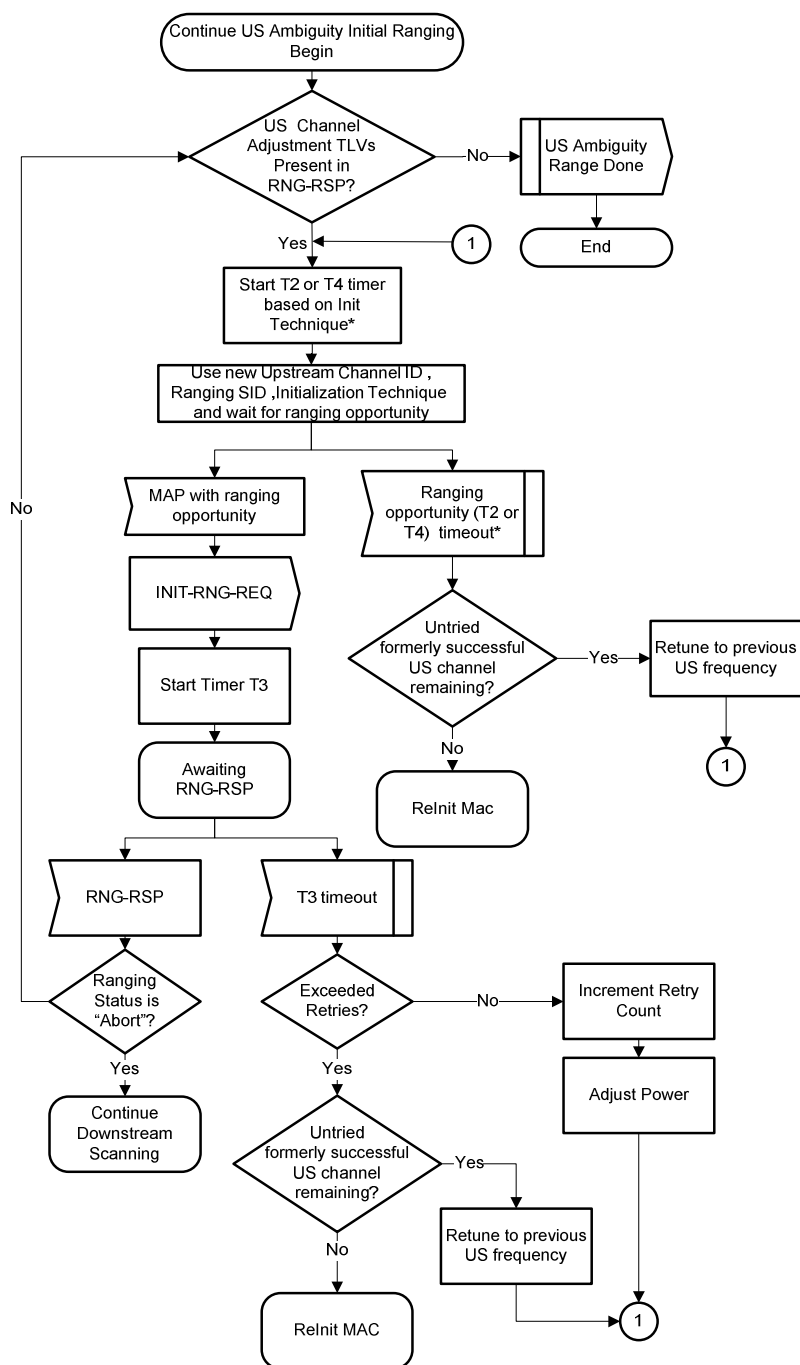
Figure 10-7: Bonded Initial Ranging

Once a candidate upstream channel has been chosen for upstream ambiguity resolution, the CM MUST attempt Bonded Initial Ranging as shown in figure 10-7 and as described below:

- 1) The CM MUST transmit a B-INIT-RNG-REQ message to the CMTS with the MD-DS-SG-ID that was determined in clause 10.2.3.3 if the MD-DS-SG-ID could be determined or an MD-DS-SG-ID of zero if an MD-DS-SG-ID could not be determined. The CM starts the T3 timer upon transmission of the B-INIT-RNG-REQ message and then waits for a response.
- 2) If the CM receives a RNG-RSP message with a Ranging Status other than Abort, then Bonded Initial Ranging is considered successful and the CM proceeds to the operations described in clause 10.2.3.5.2. If the CM receives a RNG-RSP message with a Ranging Status of Abort, the CM continues scanning for a new downstream channel (clause 10.2.2).
- 3) If timer T3 expires before receiving a RNG-RSP message and the retry limit has not been exceeded, then the CM MUST adjust power and return to step 1.
- 4) If timer T3 expires before receiving a RNG-RSP message and the retry limit has been exceeded, then the CM considers the bonded initial ranging process on the current UCID to have failed.

10.2.3.5.2

Continue US Ambiguity Initial Ranging



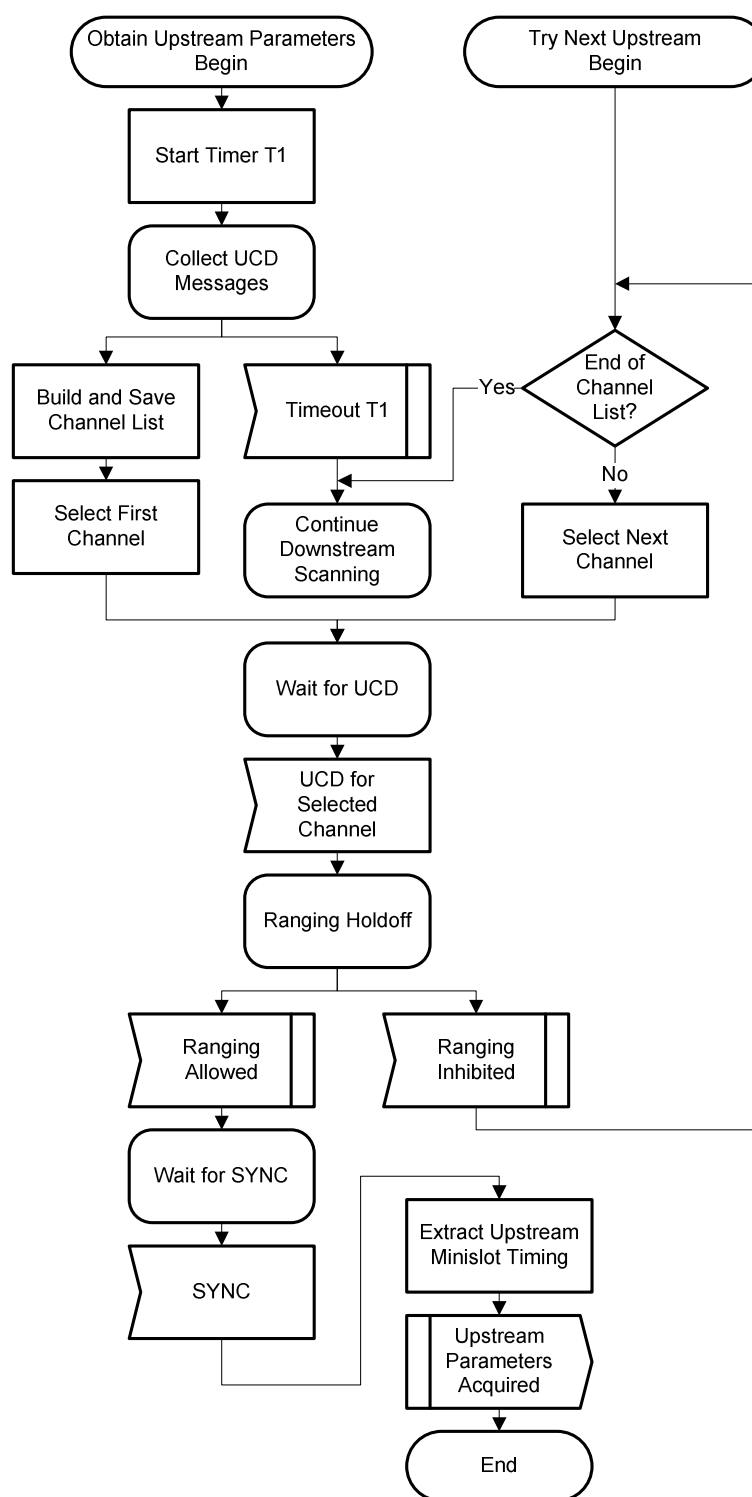
NOTE: The ranging opportunity timeout is dependent on the Initialization Technique attribute in the current adjustment request. If Technique 1 is used then the timeout value is T2. If Initialization Techniques 2 or 3 are used the timeout value is T4.

Figure 10-8: Continue US Ambiguity Initial Ranging

Once Bonded Initial Ranging has succeeded, the CM MUST continue the process of initial ranging on each channel as controlled by the CMTS as shown in figure 10-8:

- 1) If the RNG-RSP message received during Bonded Initial Ranging (see clause 10.2.3.5.1) contains Upstream Channel Adjustment TLVs, then the CM uses the Upstream Channel Adjustment TLVs and performs initial ranging (INIT-RNG-REQ) using the new Upstream Channel ID and corresponding UCD, Temp SID (if present) and Initialization Technique. To speed up the ranging process, additional ranging parameter offsets may also be included. The CMTS may respond to successive ranging request messages with a series of RNG-RSP messages containing different Upstream Channel Adjustment TLVs as it attempts to assign a MD-US-SG-ID to the CM.
- 2) If any Upstream Channel Adjustment is unsuccessful, then the CM tries to use initial ranging on an upstream channel that the CM had previously successfully ranged upon. The act of initial ranging on a previous channel tells the CMTS that the Upstream Channel Adjustment was unsuccessful.
- 3) If initial ranging on that previous channel is no longer successful, then the CM tries initial ranging on any other previously successful upstream channel. When all previously successful upstream channels have been tried without success, the CM reinitializes the MAC with a CM Initialization Reason of ALL_US_FAILED.

10.2.3.6 Obtain Upstream Parameters / Try Next Upstream (DOCSIS 2.0 Initialization)

**Figure 10-9: Obtain Upstream Parameters**

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 (see figure 10-9).

UCD messages are transmitted periodically from the CMTS for all available upstream channels and are addressed to the all CM MAC multicast address. The CM MUST determine whether it can use the upstream channel from the channel description parameters.

The CM MUST collect all UCDs with different channel ID fields to build a set of usable channel IDs. If no channel can be found after a suitable timeout period, the CM MUST continue scanning to find another downstream channel.

The CM MUST determine whether it can use the upstream channel from the channel description parameters. If the channel is not suitable, the CM MUST try other channels until it finds a usable channel.

Before attempting initial ranging on an upstream, the CM categorizes the available upstreams into the following types based on the UCD for each channel:

- Type 1, with a UCD (MAC management message type 2) offering DOCSIS 1.x burst descriptors only.
- Type 2, with a UCD (MAC management message type 2) offering both DOCSIS 2.0 TDMA and DOCSIS 1.x burst descriptors.
- Type 3, with a UCD (MAC management message Type 29 with no Type 35 present) for a DOCSIS 2.0 Upstream.
- Type 4, with a UCD (MAC management message Type 35 present and optionally Type 29 present also) for a DOCSIS 3.0 Upstream.

The CM MUST have non-volatile storage in which channel ID of the last upstream on which the CM successfully completed registration is stored. If multiple upstreams are available, the CM MUST attempt to use the one that matches this stored channel ID. If none of the available upstreams match that stored ID or if the CM is unable to successfully complete initial ranging on the matching channel, then the CM MUST preferentially select upstream channels in the following order: Type 4 channels are first, followed by Type 3 channels, followed by Type 2 channels and Type 1 channels are last. The CM MUST NOT begin initial ranging on a lower type number upstream until it has allowed sufficient time, at least the UCD Interval (refer to annex B), to determine if a higher type upstream is available. If initial ranging fails on a Type 4 upstream, the CM MUST ensure that it has allowed sufficient time to detect any other Type 4 upstreams that are available before moving on to a Type 3, Type 2 or Type 1 upstream. Of course, once the CM has waited enough time to ensure that it knows about any available Type 4 upstreams, it will also know about any available Type 3 upstreams and it MUST try them in preference to any Type 2 or Type 1 upstreams.

If the channel is suitable, the CM MUST extract the parameters for this upstream from the UCD. It then performs Ranging Holdoff and waits for the next SYNC message and extracts the upstream mini-slot timestamp from this message.

10.2.3.6.1 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 annex B. These messages are addressed to all CMs. Refer to figure 10-10.

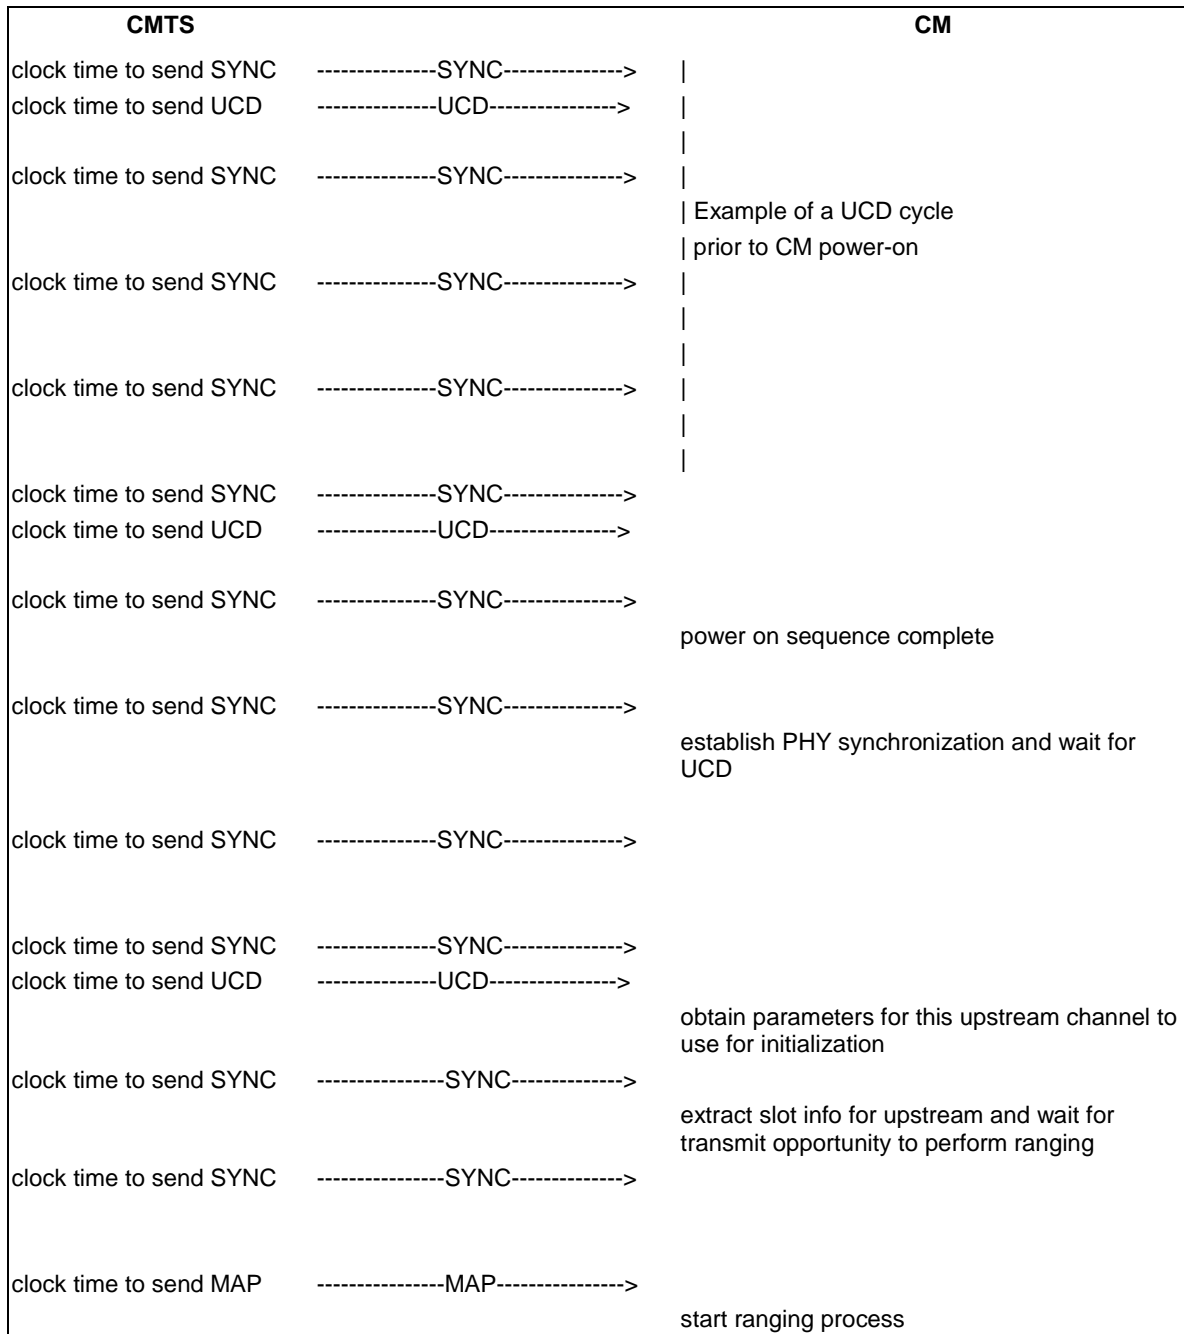
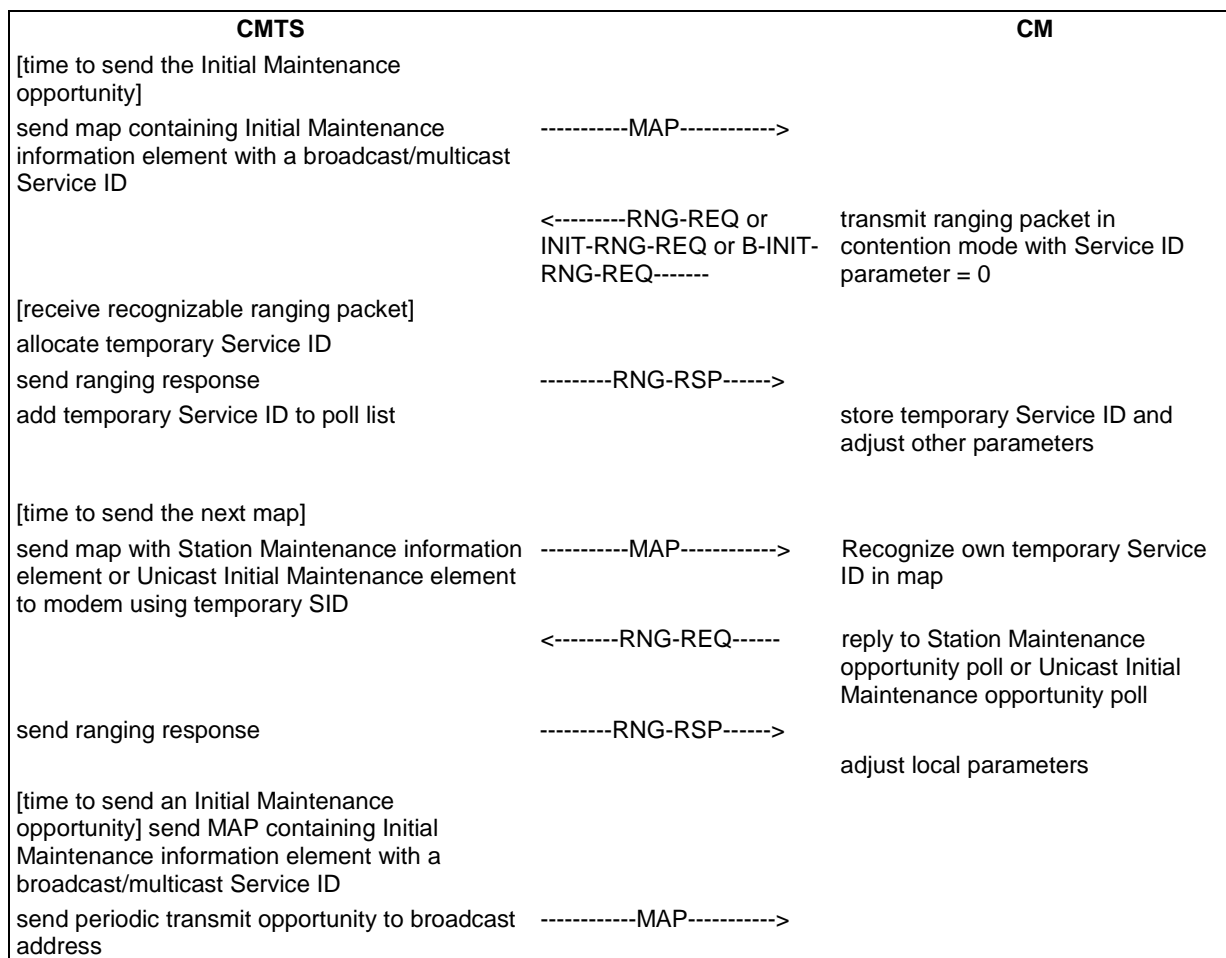


Figure 10-10: Message Flows During Scanning and Upstream Parameter Acquisition

10.2.3.7 Ranging and Automatic Adjustments

The ranging and adjustment process is fully defined in clause 6 and in the following clauses. 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 10-11 through figure 10-13.

NOTE: MAPs are transmitted as described in clause 6.



NOTE: The CMTS MUST allow the CM at least the CM Ranging Response time (annex B) to process the previous RNG-RSP (i.e. to modify the transmitter parameters) before sending the CM a unicast ranging opportunity.

Figure 10-11: Ranging and Automatic Adjustments Procedure

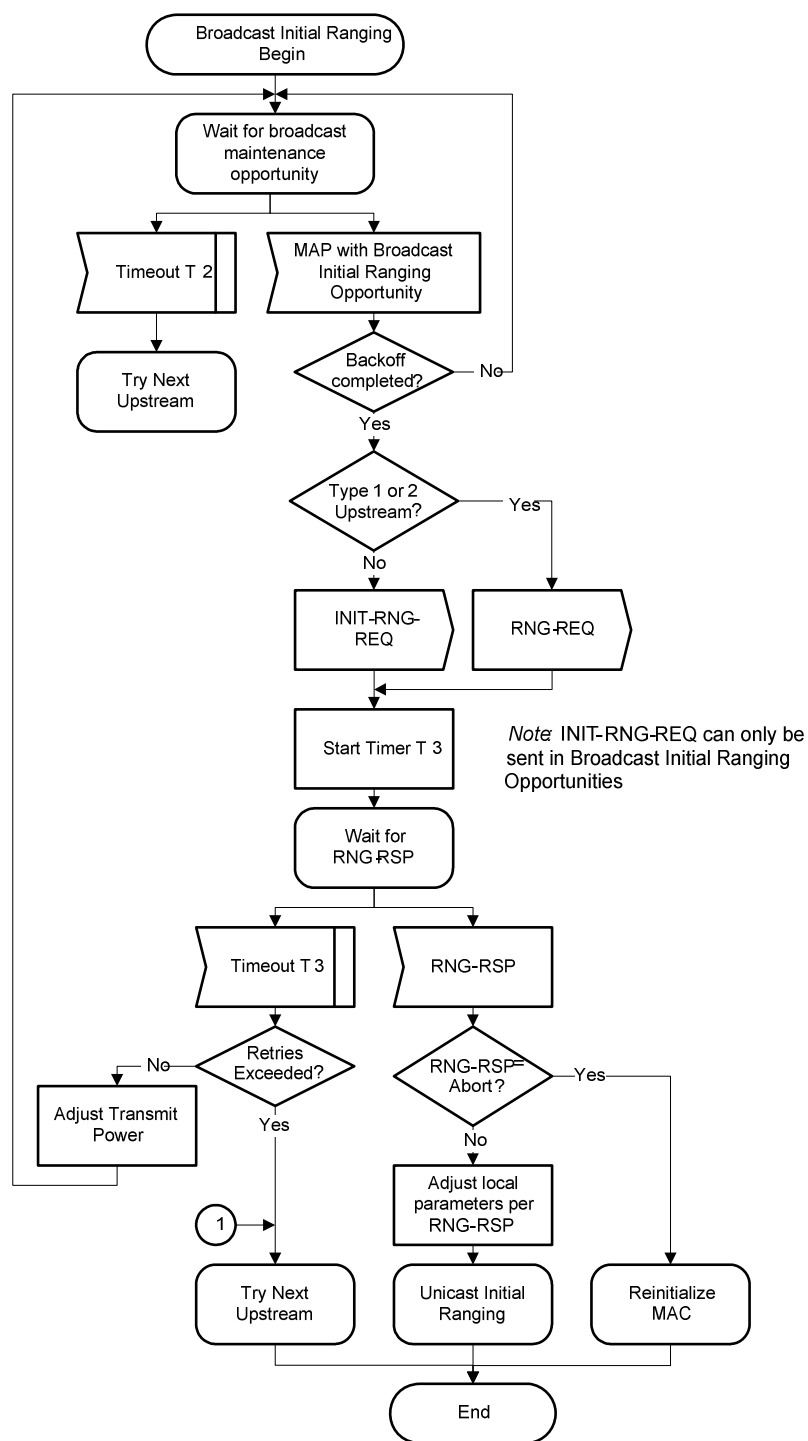


Figure 10-12: Broadcast Initial Ranging - CM

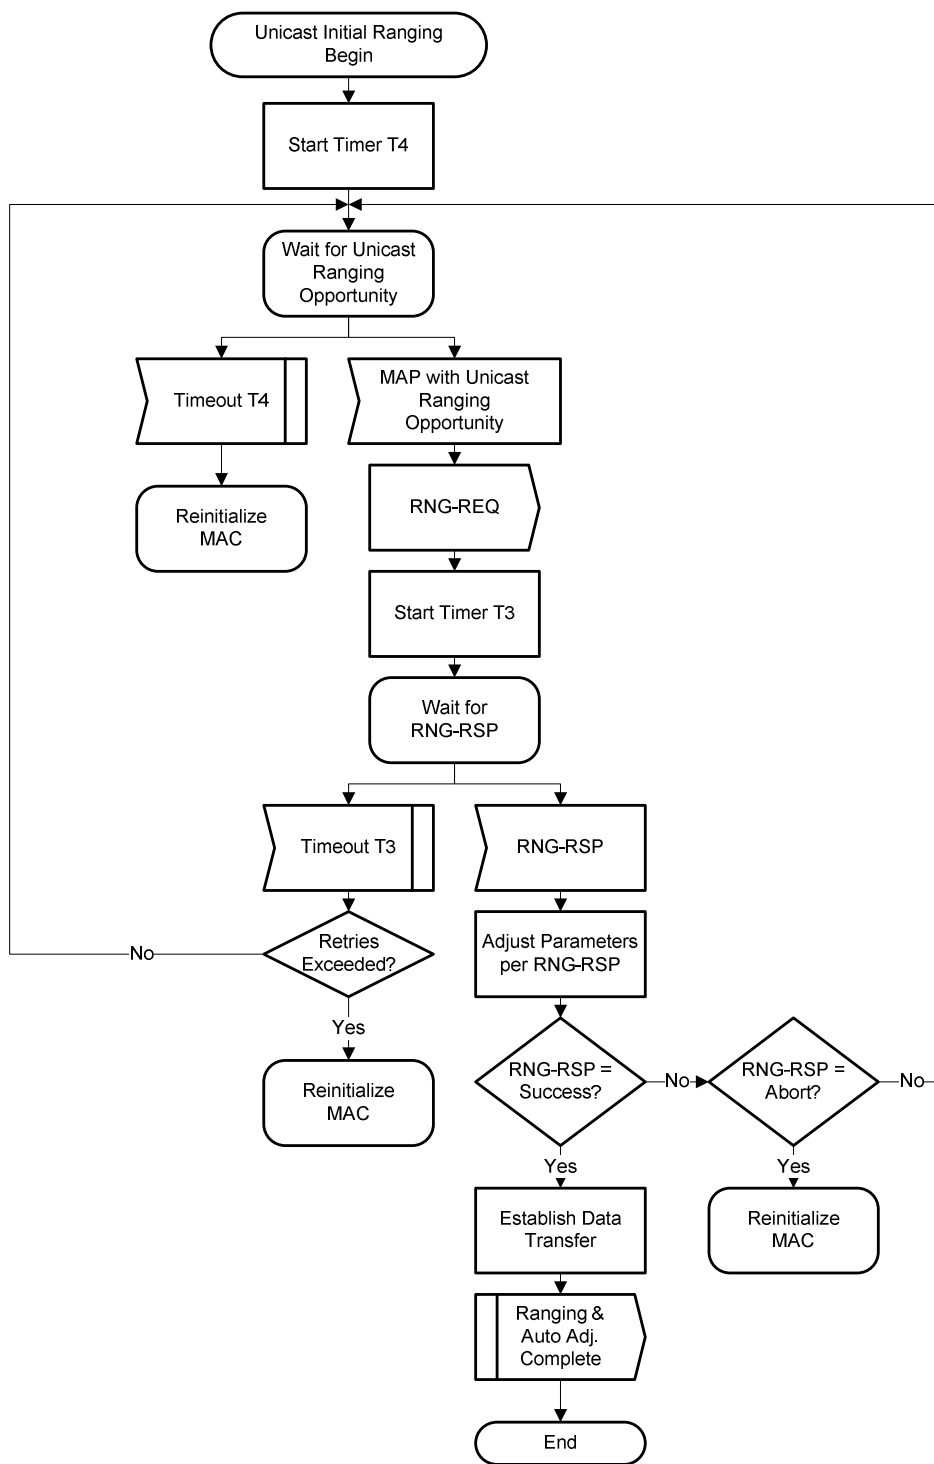


Figure 10-13: Unicast Initial Ranging - CM

10.2.3.7.1 Adjust Transmit Parameters

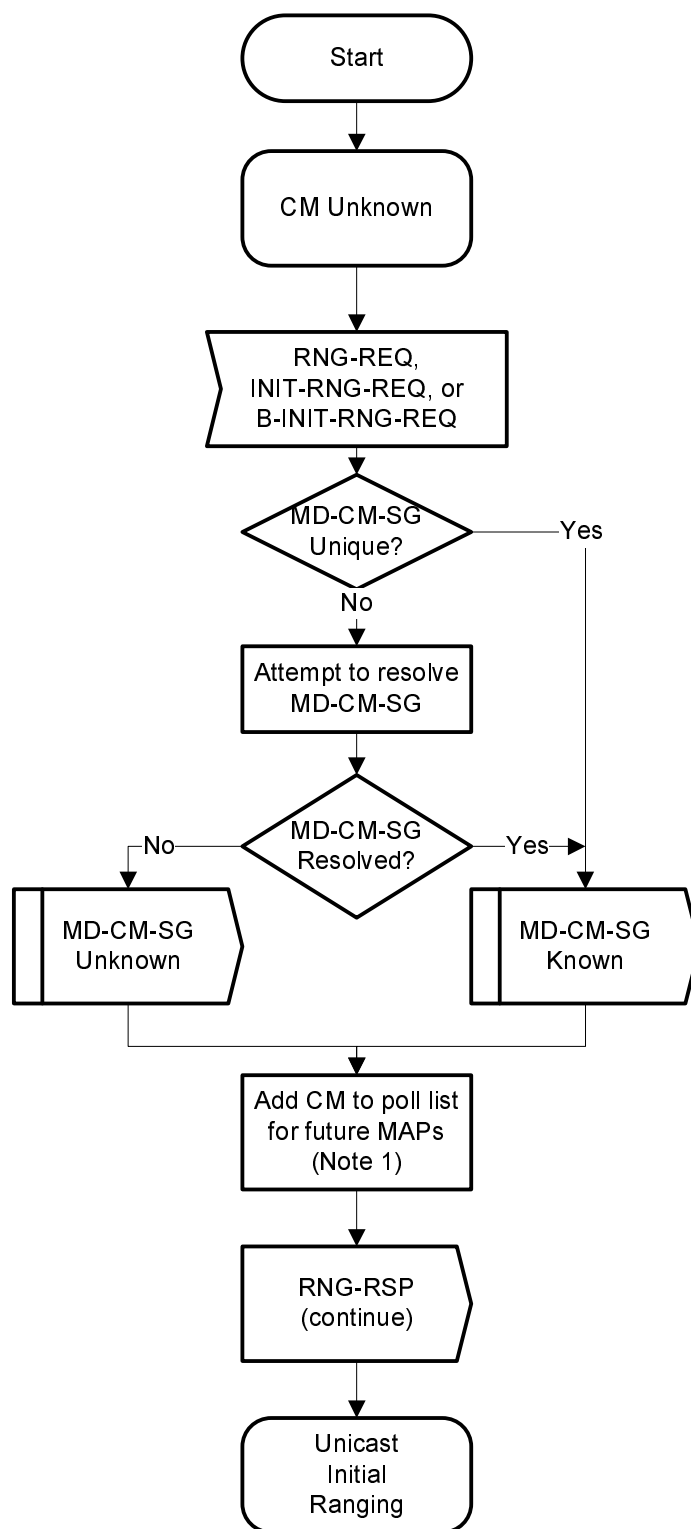
Upon receipt of a RNG-RSP message, the CM MUST reduce or increase the power by the specified amount in RNG-RSP messages.

Adjustment of local parameters (e.g. transmit power) in a CM as a result of the non-receipt of a RNG-RSP is considered to be implementation-dependent with the following restrictions (refer to clause 6.4.6):

- The CM MUST ensure that all transmit parameters are within the approved range at all times.
- For ranging prior to starting registration, the CM MUST start power adjustment from the minimum value unless a valid power is available from non-volatile storage for the upstream channel. If a valid power level for the upstream channel is available from non-volatile storage then the CM MUST use this value as a starting point.
- If Channels are being added to the TCS of modem operating in Multiple Transmit Channel Mode and no power level has been provided in the TCC encodings, then the CM MUST start power adjustment from the minimum value allowed by the Dynamic Range Window, unless a valid power is available from non-volatile storage for an upstream channel. If a valid power level for an upstream channel is available from non-volatile storage, then the CM MUST use this value as a starting point. A power level stored in non-volatile storage for the upstream channel is considered to be valid if it lies within the Dynamic Range Window.
- During initialization, prior to starting registration, the CM MUST cover its entire dynamic range within 16 retries leaving no power interval greater than 12 dB untried.
- During initial ranging on channels being added by TCC encodings the CM MUST cover the entire Dynamic Range Window within 16 retries, leaving no power interval greater than 6dB untried.

10.2.3.8 CMTS Determination of Cable Modem Service Group and Initial Ranging

The CMTS MUST attempt to determine the CM-SG of an initially ranging CM and complete Initial Ranging for that CM according to figures 10-14 and 10-15 that follow.



NOTE: The poll list is a list of CMs that are currently performing unicast initial ranging.

Figure 10-14: CM-SG Determination - CMTS

The CMTS SHOULD provide CMs in the poll list frequent unicast ranging opportunities. If RNG-REQ pending-till-complete was nonzero, the CMTS SHOULD hold off the station maintenance opportunity accordingly unless needed, for example, to adjust the CM's power level. If opportunities are offered prior to the pending-till-complete expiry, the CMTS MUST NOT use the "good-enough" test which follows receipt of a RNG-REQ in figure 10-15 to judge the CM's transmit equalization until pending-till-complete expires.

A CMTS is said to consider a CM to be "known" when it provides unicast ranging opportunities to the CM. The CMTS initially considers a CM's MAC address to be unknown, represented by the "CM unknown" state of figure 10-14. While in the "CM unknown" state, upon the receipt of a B-INIT-RNG-REQ, an INIT-RNG-REQ or a RNG-REQ where the DCID, UCID and MD-DS-SG-ID (if present and non-zero) are associated with one and only one MD-CM-SG, the MD-CM-SG is considered to be "Unique" and thus "Known" by the CMTS. It is CMTS vendor-specific whether or to what degree, the CMTS attempts to determine the MD-CM-SG of CMs for which the MD-CM-SG is not "Unique" based the information available in the B-INIT-RNG-REQ, INIT-RNG-REQ or RNG-REQ, i.e. the SDL procedure named "Attempt to resolve MD-CM-SG" is vendor-specific. If the CMTS does not support such MD-CM-SG determination or cannot determine the MD-CM-SG of a CM, it considers the MD-CM-SG to be "unknown" for the CM.

The "Attempt to resolve MD-CM-SG" procedure might proceed as follows. The MD-DS-SG of the ranging CM might be uniquely determined by the MD-DS-SG-ID in the B-INIT-RNG-REQ, by the Downstream Channel ID (DCID) in the INIT-RNG-REQ or RNG-REQ message or by the particular Upstream Channel from which the B-INIT-RNG-REQ, INIT-RNG-REQ or RNG-REQ was received. If the MD-DS-SG has not been uniquely determined, the CMTS can send a RNG-RSP to the CM to override the downstream frequency to one for which the CMTS can reduce the set of possible MD-DS-SGs for the CM. At that point, if the CMTS receives a B-INIT-RNG-REQ, an INIT-RNG-REQ or a RNG-REQ message from the CM, it notes the MD-DS-SG-ID and/or DCID reported in the new B-INIT-RNG-REQ, INIT-RNG-REQ or RNG-REQ and continues checking whether the MD-DS-SG is unique. Note that if the CMTS receives a B-INIT-RNG-REQ, an INIT-RNG-REQ or a RNG-REQ with a DCID corresponding to a downstream frequency other than the requested override frequency, it indicates either that the CM was unable to detect an acceptable downstream channel at that frequency or that the CM was reset from a power-on condition. To avoid this ambiguity, a Cable Operator can implement a downstream RF topology where each CM is reached by a valid DOCSIS downstream channel at every frequency used by any DOCSIS downstream channel in the MAC Domain.

Once the MD-DS-SG is uniquely determined, the CMTS can proceed to check if the MAC Domain Upstream Service Group (MD-US-SG) is also unique.

If the combination of MD-DS-SG and the particular Upstream Channel from which the B-INIT-RNG-REQ/INIT-RNG-REQ/RNG-REQ was received does not determine the MD-US-SG, the CMTS can send a RNG-RSP to continue ranging and use the Upstream Channel Adjustment TLV to override the Upstream Channel ID (UCID). In the RNG-RSP to a CM which sent a B-INIT-RNG-REQ, the CMTS MAY also use the Upstream Channel Adjustment TLV specify the initialization technique for the CM to use on the overridden UCID. At that point, if the CMTS receives an INIT-RNG-REQ or RNG-REQ from an upstream channel on the frequency of the overridden UCID, the CMTS adds the UCID of the actual upstream channel from which the INIT-RNG-REQ or RNG-REQ was received to a known set of Upstream Channels reaching the CM and continues checking whether the MD-US-SG is uniquely determined. If the CMTS receives an INIT-RNG-REQ/RNG-REQ from a different frequency than the overridden UCID, it indicates that the CM was unable to range on the overridden UCID's frequency. One possibility is that the CM is in an MD-US-SG that omits any Upstream Channel on the overridden UCID's frequency.

While performing a RNG-RSP downstream frequency override or RNG-RSP Upstream Channel Adjustment override, if the CMTS receives no ranging request, it can remove the CM from its set of known CMs.

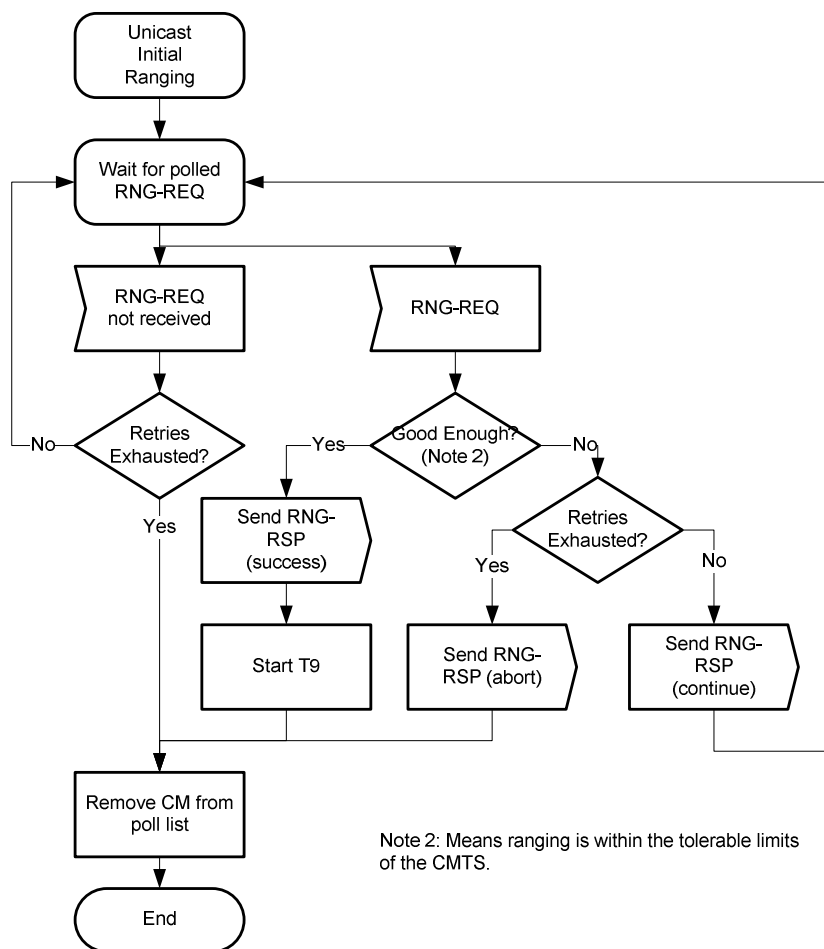


Figure 10-15: Unicast Initial Ranging - CMTS

10.2.4 Authentication

Once a CM has completed ranging, if Early Authentication and Encryption (EAE) is enabled in the MDD the CM will initiate EAE before continuing with the initialization process. EAE helps prevent unauthorized CMs from accessing IP provisioning servers and provides confidentiality/privacy for IP provisioning messages between the CM and CMTS. See [15] for details.

10.2.5 Establishing IP Connectivity

The CM performs IP provisioning in one of four modes: IPv4 Only, IPv6 Only, Alternate Provisioning Mode (APM) and Dual-stack Provisioning Mode (DPM). The CM determines the IP provisioning mode via the CmMdcfg management object defined in [10].

If the management object is set to 'honor MDD', the default setting, the CM determines the IP provisioning model by the absence of the MDD message or by the TLVs in the MDD message (see clause 6.4.28). If the CM does not receive an MDD from the CMTS, the CM MUST perform IPv4 Only provisioning. If the CM receives an MDD from the CMTS, the CM MUST use the provisioning mode directed by the MDD IP Provisioning Mode TLV.

As shown in figure 10-1, the IP provisioning process begins after the completion of ranging or EAE if enabled and ends with an IP provisioning success or failure, i.e. with the CM in either IP Connectivity Successful or IP Connectivity Failed state. As shown in figure 10-1, if the CM finishes IP provisioning successfully, it proceeds with registration; and if it does not, it continues scanning for a new downstream channel.

The Cable Modem performing IP provisioning MUST follow the operational flow of figure 10-16 through figure 10-22 to arrive at an IP Connectivity Successful or IP Connectivity Failed state. Figure 10-16 shows the selection of the provisioning modes. Figure 10-17 through figure 10-20 show the steps the CM takes in each of the provisioning modes. Figure 10-21 shows the steps the CM takes to obtain time and a configuration file. Figure 10-22 shows the process the CM follows for acquiring an IPv6 address. The acquisition of an IPv4 address, done through DHCPv4, is shown as part of figures 10-17, 10-19 and 10-20.

Once the CM is registered, any applications and services running on the CM, such as SNMP, use the IP version (v4 or v6) through which the CM obtains the configuration file used for registration, unless the CM is directed to use DPM. When the CM uses DPM, the applications and services running on the CM use both IP versions.

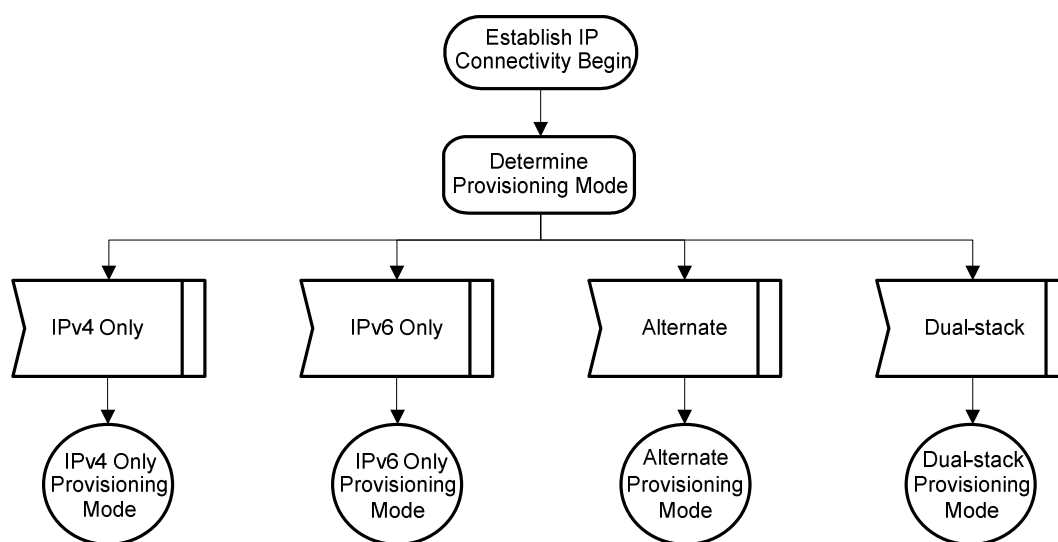
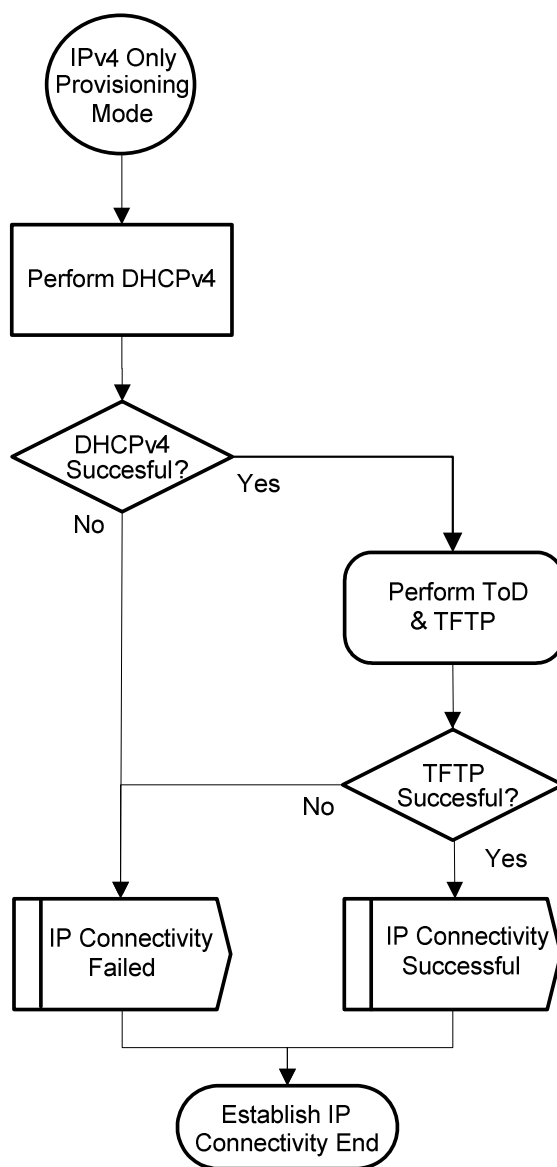


Figure 10-16: Establish IP Connectivity

**Figure 10-17: IPv4 Only Provisioning Mode**

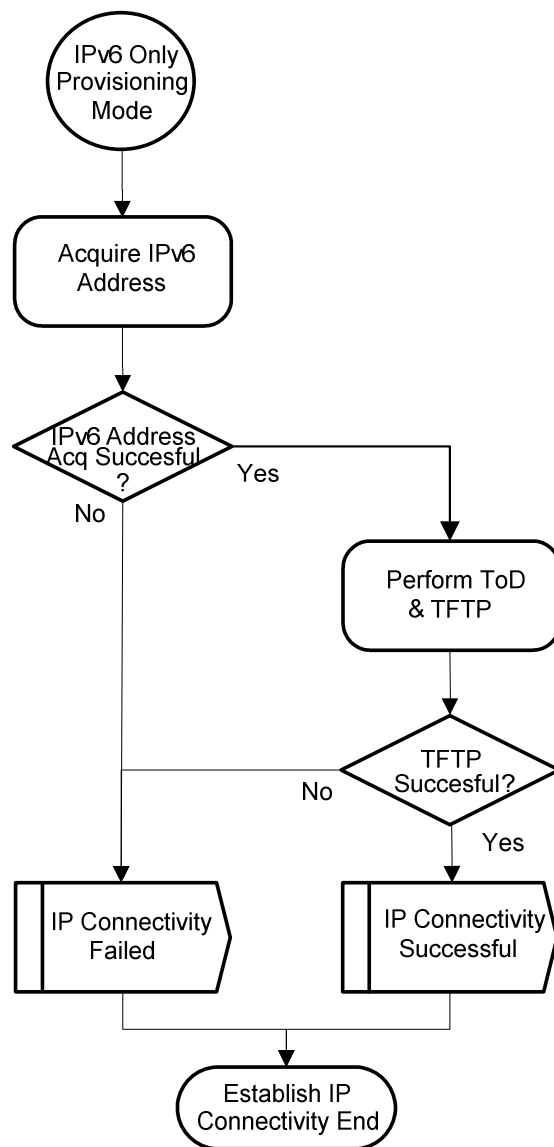
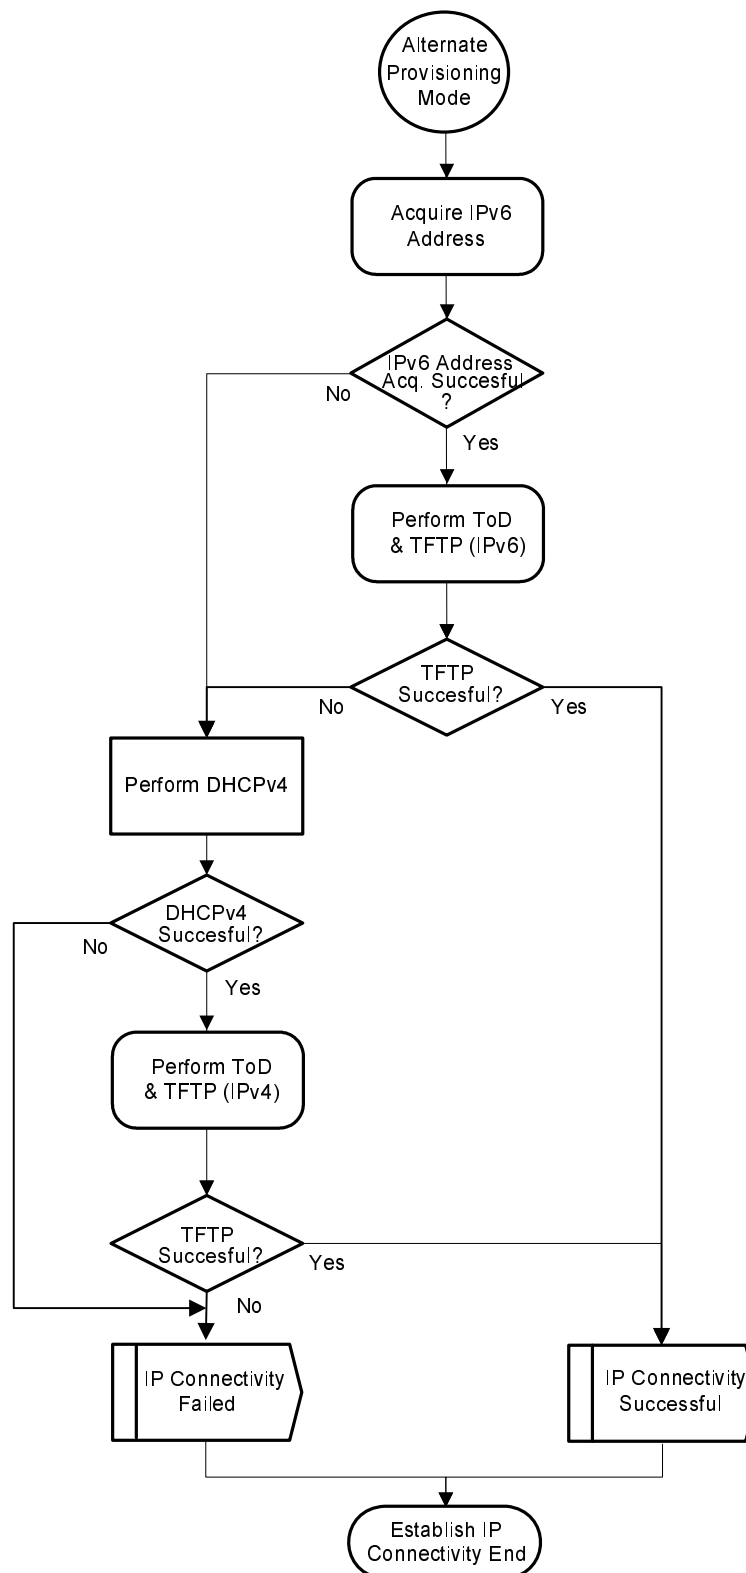
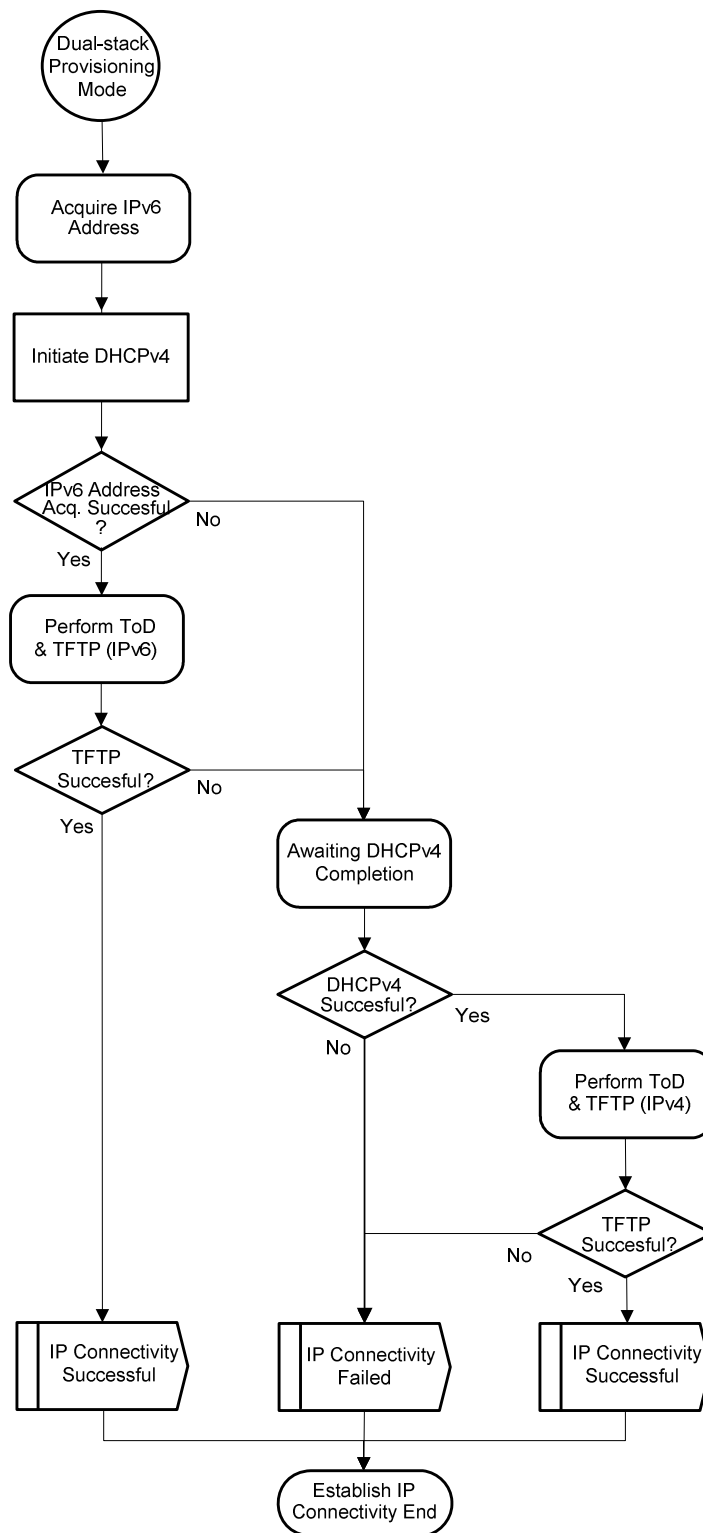


Figure 10-18: IPv6 Only Provisioning Mode

**Figure 10-19: Alternate Provisioning Mode**

**Figure 10-20: Dual-stack Provisioning Mode**

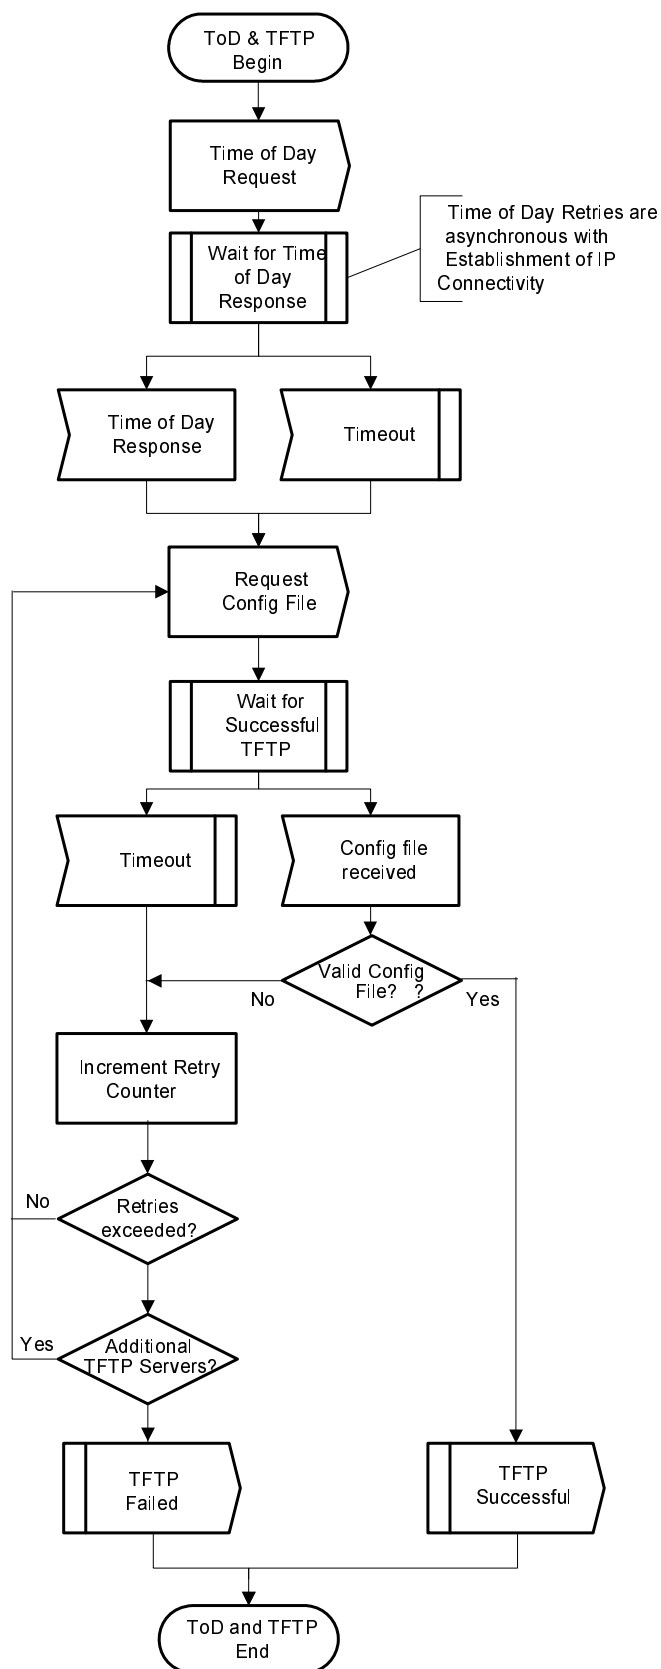


Figure 10-21: ToD and TFTP

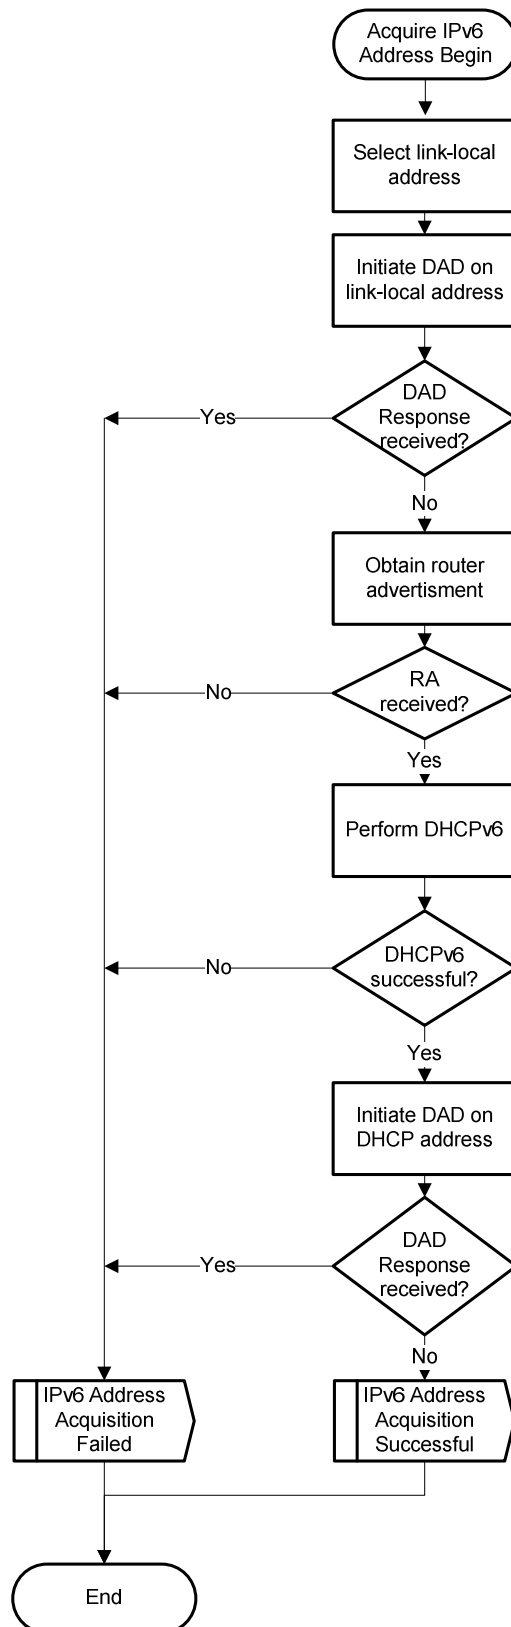


Figure 10-22: IPv6 Address Acquisition

10.2.5.1 Establish IPv4 Network Connectivity

This clause describes how the CM is provisioned with an IPv4 address and associated parameters. The requirements in this clause apply to CMs using IPv4 provisioning. A CM uses IPv4 provisioning when it does not receive an MDD from the CMTS, when the MDD indicates IPv4 Only provisioning or DPM or when the MDD indicates APM and IPv6 provisioning fails.

The CM **MUST** use DHCPv4 [37] in order to obtain an IP address and other parameters needed to establish IP connectivity in the following cases:

- The MDD is not present.
- The MDD is present and it indicates IPv4 Only provisioning.
- The MDD is present and it indicates DPM.
- The MDD is present and it indicates APM and IPv6 provisioning fails.

Figure 10-23 shows the DHCPv4 message sequence.

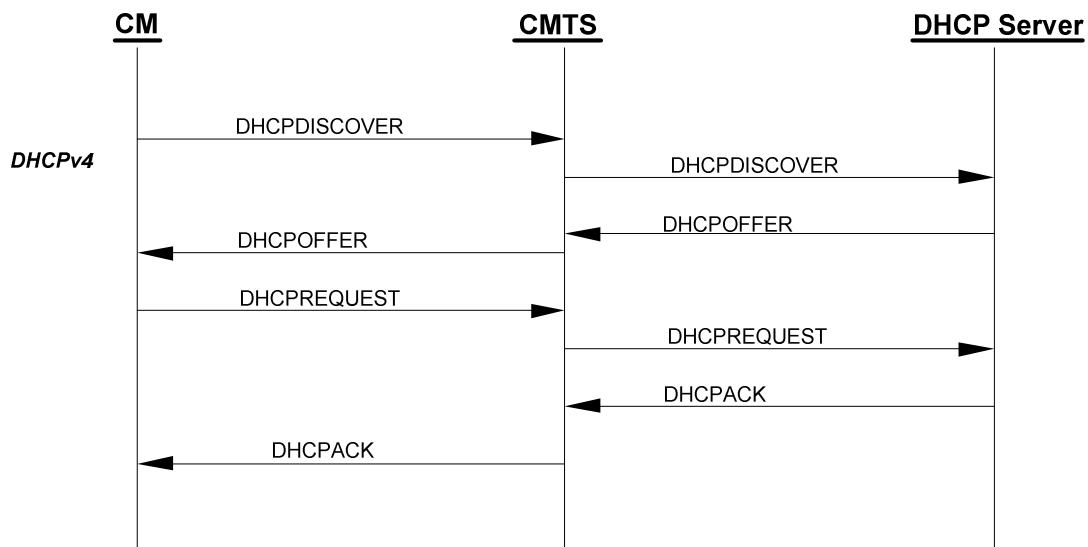


Figure 10-23: IPv4 Provisioning Message Flow

The CM may receive multiple DHCPOFFER messages in response to its DHCPDISCOVER message. If a received DHCPOFFER message does not include all of the required DHCPv4 fields and options as described in clause 10.2.5.1.1, the CM **MUST** discard the DHCPOFFER message and wait for another DHCPOFFER message. If none of the received DHCPOFFER messages contain all the required DHCPv4 fields and options, the CM retransmits the DHCPDISCOVER message.

The backoff values for retransmission of DHCPDISCOVER messages **SHOULD** be chosen according to a uniform distribution between the minimum and maximum values in the rows of table 10-1.

Table 10-1: DHCP Backoff Distribution Values

Backoff Number	Minimum (s)	Maximum (s)
1	3	5
2	7	9
3	15	17
4	31	33
5	63	65

The CM **SHOULD** also implement a different retransmission strategy for the RENEWING and REBINDING states, as recommended in [37], which is based on one-half of the remaining lease time.

The CM MUST limit the number of retransmissions to five or fewer for the DHCPDISCOVER and DHCPREQUEST messages.

[49] describes an extension to DHCPv4 that allows a DHCP server to send a FORCERENEW message that forces a client to renew its lease. The CM MUST ignore all received FORCERENEW messages.

10.2.5.1.1 DHCPv4 Fields Used by the CM

The following fields MUST be present in the DHCPDISCOVER and DHCPREQUEST messages from the CM:

- The hardware type (htype) MUST be set to 1.
- 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, using the format defined in [55].
- The parameter request list option MUST be included. The following option codes (defined in [38] and [55]) MUST be included in the list:
 - 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).
 - Option code 125 (DHCPv4 Vendor-Identifying Vendor Specific Information Option).
- Option code 60 (Vendor Class Identifier) - the following ASCII-encoded string MUST be present in Option code 60: DOCSIS 3.0:
- Option code 125 (DHCPv4 Vendor-Identifying Vendor Specific Information Options for DOCSIS 3.0 defined in [1] - and include the following sub-options in there:
 - 1) Sub-option code 1, the DHCPv4 Option Request option. The following option code MUST be included in the DHCPv4 Option Request option.
 - 2) Sub-option code 2, DHCPv4 TFTP Servers Option.
 - 3) Sub-option code 5, Modem Capabilities Encoding for DHCPv4.

The following fields are expected in the DHCP OFFER and DHCPACK messages returned to the CM. The CM MUST configure itself with the listed fields from the DHCPACK:

- The IP address to be used by the CM (yiaddr) (critical).
- The IP addresses of the TFTP servers for use in the next phase of the boot process (DHCPv4 TFTP Servers option defined in [1] or siaddr) (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 UTC (Time Offset, option 2). This is used by the CM to calculate a time for use in error logs (non-critical).
- A list of addresses of one or more routers to be used for forwarding IP traffic originating from the CM's IP stack (Router Option, option 3). The CM is not required to use more than one router IP address for forwarding (non-critical).

- A list of ToD 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 [9] (non-critical).

If a critical field is missing or invalid in the DHCPACK received during initialization, the CM MUST:

- 1) Log an error.
- 2) Proceed as if the acquisition of the IPv4 address through DHCPv4 has failed; reference figures 10-17, 10-19 and 10-20.

If a non-critical field is missing or invalid in the DHCPACK received during initialization, the CM MUST log a warning, ignore the field and continue the IPv4 provisioning process.

If the yiaddr field is missing or invalid in the DHCPACK received during a renew or rebind operation, the CM MUST log an error and reinitialize its MAC with a CM Initialization Reason of BAD_DHCP_ACK.

If any other critical or non-critical field is missing or is invalid in the DHCPACK received during a renew or rebind operation, the CM MUST log a warning, ignore the field if it is invalid and remain operational.

10.2.5.1.2 Use of T1 and T2 Timers

The CM MUST initiate the lease renewal process when timer DHCP-T1 expires. The CM MUST initiate the lease rebinding process when timer DHCP-T2 expires. Timers DHCP-T1 and DHCP-T2 are called T1 and T2, respectively, in the DHCP specifications. If the DHCP server sends a value for DHCP-T1 to the CM in a DHCP message option, the CM MUST use that value. If the DHCP server does not send a value for DHCP-T1, the CM MUST set DHCP-T1 to one half of the duration of the lease [37]. If the DHCP server sends a value for DHCP-T2 to the CM in a DHCP message option, the CM MUST use that value. If the DHCP server does not send a value for DHCP-T2, the CM MUST set DHCP-T2 to seven-eighths of the duration of the lease [37].

10.2.5.1.2.1 DHCPv4 Renew Fields Used by the CM

It is possible during the DHCPv4 renew operation that the CM will receive updated fields in the DHCPACK message.

If any of the IP address (yiaddr), the Subnet Mask or the Next Hop Router (router option) are different in the DHCPACK than the current values used by the CM, the CM MUST follow one of the following two:

- Change to using the new values without reinitializing the CM; or
- reinitialize MAC.

If the Config File Name or the SYSLOG server address are different in the DHCPACK than the current values used by the CM, the CM MUST ignore the new fields.

If the Time Offset value is different in the DHCPACK than the current value used by the CM, the CM MUST update the internal representation of time based on the new Time Offset value.

If the Time server address is different in the DHCPACK than the current value used by the CM, the CM MUST update the time server address with the new value. This will cause the CM to use the new address(es) on future ToD requests (if any).

10.2.5.1.3 CMTS Requirements

In order to assist the DHCP server in differentiating between a DHCPDISCOVER sent from a CM and a DHCPDISCOVER sent from a CPE:

- The CMTS DHCPv4 relay agent MUST support the DHCP Relay Agent Information Option (RAIO) [48]. Specifically, the CMTS DHCPv4 relay agent MUST add an RAIO to the DHCPDISCOVER message before relaying the message to a DHCP server. The RAIO MUST include the 48 bit MAC address of the RF-side interface of the CM generating or bridging the DHCPDISCOVER in the agent remote ID sub-option field [48].

- If the CMTS is a router, the CMTS DHCPv4 relay agent **MUST** use a giaddr field to differentiate between CM and CPE clients if they are to be provisioned in different IP subnets. The DHCPv4 relay agent in a bridging CMTS **SHOULD** provide this function.
- DHCPv4 Relay Agent CMTS Capabilities option, containing the value "3.0" for the CMTS DOCSIS Version Number [1].

The CMTS DHCPv4 relay agent **MAY** support the DHCP Relay Agent Service Class Information sub option [1].

10.2.5.2 Establish IPv6 Network Connectivity

This clause describes how the CM is provisioned with an IPv6 address and associated configuration parameters. The requirements in this clause apply only to CMs instructed to use IPv6 provisioning. A CM uses IPv6 provisioning when the MDD indicates IPv6 Only provisioning, DPM or APM.

Figure 10-24 illustrates the message flows in IPv6 provisioning.

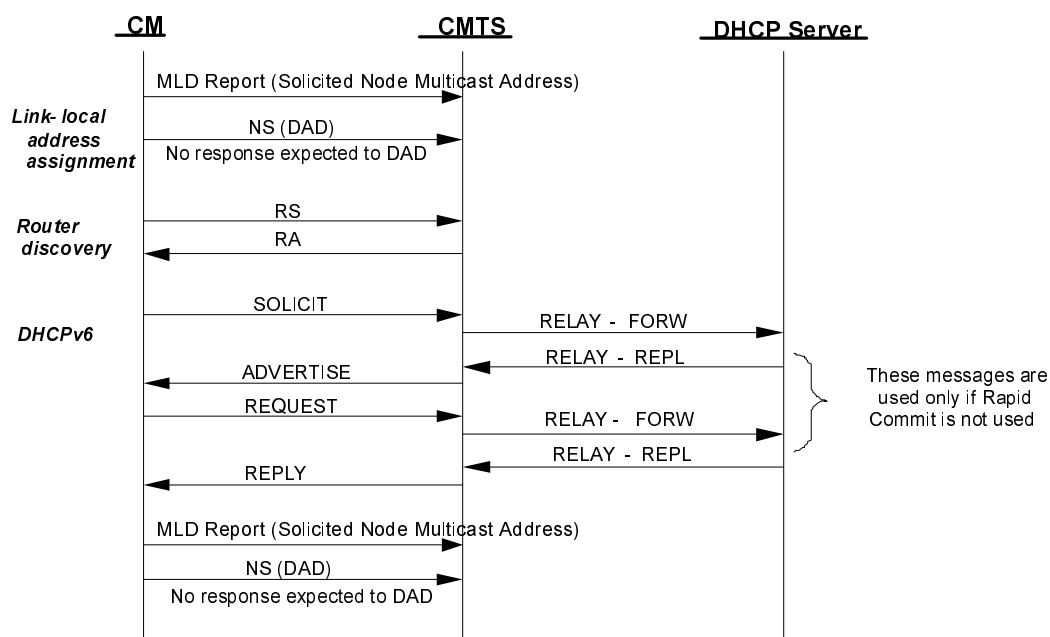


Figure 10-24: IPv6 Provisioning Message Flow

The CM establishes IPv6 connectivity including assignment of:

- Link-local address.
- Default router.
- IPv6 management address.
- Other IPv6 configuration.

These steps are described in the following clauses.

10.2.5.2.1 Obtain Link-Local Address

The CM MUST construct a link-local address for its management interface according to the procedure in [60]. The CM MUST use the EUI-64 (64-bit Extended Unique Identifier) as a link-local address for its management interface as described in [53]. The CM management interface MUST join the all-nodes multicast address and the solicited-node multicast address of the constructed link-local address [60]. When joining the solicited-node multicast address of the constructed link-local address, the CM MUST immediately report this address in an unsolicited MLD Report; the all-nodes multicast address is not reported [46]. The CM MUST use Duplicate Address Detection (DAD), as described in [60], to confirm that the constructed link-local address is not already in use. If the CM determines that the constructed link-local address is already in use, the CM MUST report the event in its local log, reinitialize its MAC with a CM Initialization Reason of LINK_LOCAL_ADDRESS_IN_USE and resume scanning to find another downstream channel. If a CM fails Duplicate Address Detection the CM MUST NOT assign the tentative EUI-64 address to the interface.

If the link-local address used by a CM as a source IPv6 address is not constructed from the CM's MAC address, the CMTS MUST report the event in its local log and deny the CM's attempt to register. A CMTS that acts as a proxy for ND MUST send a NA message in response to a NS from a CM if it detects that another CM has already assigned the target address on the link. When the CMTS sends the NA message, it MUST log the event and prevent registration by the CM with the duplicate address. A CMTS that acts as a proxy for ND MUST NOT forward any Neighbor Discovery Packets received on an upstream channel to any downstream channel.

10.2.5.2.2 Obtain default routers

The CM MUST perform router discovery as specified in [59]. The CM identifies neighboring routers and default routers from the received RAs. If the CM does not receive a properly formatted response to the Router Solicitation message within the retransmission requirements defined in [59], the CM MUST proceed as if IPv6 Address Acquisition has failed.

10.2.5.2.3 Obtain IPv6 management address and other configuration parameters

A CM MUST examine the contents of RAs that it receives. The CM obeys the following rules:

- If the M bit in the RA is set to 1, the CM MUST use DHCPv6 to obtain its management address and other configuration information (and ignore the O bit).
- If there are no prefix information options in the RA, the CM MUST NOT perform SLAAC.
- If the RA contains a prefix advertisement with the A bit set to 0, the CM MUST NOT perform SLAAC on that prefix.

The CM sends a DHCPv6 Solicit message as described in [51]. The Solicit message MUST include:

- A Client Identifier option containing the DUID (DHCP Unique Identifier) for this CM as specified by [51]. The CM can choose any one of the rules to construct the DUID according to clause 9.1 of [51].
- An IA_NA (Identity Association for Non-temporary Addresses) option to obtain its IPv6 management address.
- A Vendor Class option containing 32-bit number 4491 (the Cable Television Laboratories, Inc. enterprise number) and the string "docsis 3.0".
- A Vendor-specific option containing:
 - 1) TLV5 option (reference [1]) containing the encoded TLV5s describing the capabilities of CM information option in clause C.1.3.1.
 - 2) Device ID option containing the MAC address of the HFC interface of the CM.
 - 3) ORO option requesting the following vendor specific options:
 - a) Time Protocol Servers.
 - b) Time Offset.
 - c) TFTP Server Addresses.

- d) Configuration File Name.
 - e) SYSLOG Server Addresses.
- A Rapid Commit option indicating that the CM is willing to perform a 2-message DHCPv6 message exchange with the server.

The CM MUST use the following values for retransmission of the Solicit message (see [51] for details):

- IRT (Initial Retransmission Time) = SOL_TIMEOUT.
- MRT (Maximum Retransmission Time) = SOL_MAX_RT.
- MRC (Maximum Retransmission Count) = 4.
- MRD (Maximum Retransmission Duration) = 0.

The CM MUST use the values for retransmission of the Request message defined in [51].

If the number of retransmissions is exhausted before the CM receives an Advertise or Reply message from a DHCP server, the CM MUST consider IPv6 address acquisition to have failed.

The CM MUST support the Reconfigure Key Authentication Protocol as described in [51].

The DHCPv6 server may be configured to use a 2 message Rapid Commit sequence. The DHCP server and CM follow [51] in the optional use of the Rapid Commit message exchange.

The DHCPv6 server responds to Solicit and Request messages with Advertise and Reply messages (depending on the use of Rapid Commit). The Advertise and Reply messages may include other configuration parameters, as requested by the CM or as configured by the administrator to be sent to the CM. If any of the following options is absent from the Advertise message, the CM MUST discard the message and wait for other Advertise messages. If any of the following options is absent from the Reply message, the CM MUST consider IPv6 address acquisition to have failed:

- The IA_NA option received from the CM, containing the IPv6 management address for the CM.
- A Vendor-specific Information option containing the following sub-options (refer to [1]):
 - 1) Time Protocol Servers option.
 - 2) Time Offset option.
 - 3) TFTP Server Addresses option.
 - 4) Configuration File Name option.
 - 5) Syslog Server Addresses option.

The CM management interface MUST join the all-nodes multicast address and the solicited-node multicast address of the IPv6 address acquired through DHCPv6 [60]. When joining the solicited-node multicast address of the IPv6 address acquired through DHCPv6, the CM MUST immediately report this address in an unsolicited MLD Report; the all-nodes multicast address is not reported [46]. The CM MUST perform a Duplicate Address Detection with the IPv6 address acquired through DHCPv6. If the CM determines through DAD the IPv6 address assigned through DHCPv6 is already in use by another device, the CM MUST send a DHCP Decline message to the DHCP server indicating that it has detected that a duplicate IP address exists on the link. The CM MUST NOT continue using this IP address. The CM MUST log an error and consider the IPv6 address acquisition to have failed.

10.2.5.2.4 Use of T1 and T2 Timers

The CM MUST initiate the lease renewal process when timer DHCP-T1 expires. The CM MUST initiate the lease rebinding process when timer T2 expires. Timers DHCP-T1 and DHCP-T2 are called T1 and T2, respectively, in the DHCP specifications. If the DHCP server sends a value for DHCP-T1 to the CM in a DHCP message option, the CM MUST use that value. If the DHCP server does not send a value for DHCP-T1, the CM MUST set DHCP-T1 to 0,5 of the duration of the lease [37]. If the DHCP server sends a values for DHCP-T2 to the CM in a DHCP message options, the CM MUST use that value. If the DHCP server does not send a value for DHCP-T2, the CM MUST set DHCP-T2 to 0,875 of the duration of the lease [51].

10.2.5.2.4.1 DHCPv6 Renew Fields Used by the CM

It is possible during the DHCPv6 renew operation that the CM will receive updated fields in the DHCP Reply message.

If the CM IPv6 Management Address (IA_NA option) is different in the DHCP Reply than the current value used by the CM, the CM MUST follow one of the following two:

- change to using the new IPv6 Management Address without reinitializing the CM; or
- reinitialize MAC.

If the following values, TFTP configuration file name (Vendor Specific Option), the Syslog servers (Vendor Specific Option) or the Reconfigure Accept option are different in the DHCP Reply than the current values used by the CM, the CM MUST ignore the new fields.

If the Time Offset Option value is different in the DHCP Reply than the current value used by the CM, the CM MUST update the internal representation of time based on the new Time Offset value.

If the Time Protocol Servers option in the DHCP Reply is different than the current value used by the CM, the CM MUST update the time server address with the new value. This will cause the CM to use the new address(es) on future ToD requests (if any).

During DHCPv6 Renew or Rebind, the CM MUST remain operational through changes, deletions or additions of any other options in the DHCPv6 Reply messages.

10.2.5.2.5 CMTS Requirements

The CMTS DHCPv6 relay agent MUST send the following DHCPv6 options to the DHCPv6 server, in any Relay-Forward messages used to forward messages from the CM to the DHCPv6 server:

- Interface-ID option [51].
- CMTS Capabilities option, containing the value "3.0" for the CMTS DOCSIS Version Number [1].
- CM MAC address option [1].
- Remote-ID option (RFC 4649 [61]).

The CMTS MUST set the Remote-ID option to the 48 bit MAC address of the RF-side interface of the CM generating or bridging the DHCPDISCOVER sent in the CL_Option_Device_ID sub-option field, as defined in [1].

In order to refresh its DHCP state information, the CMTS SHOULD support Bulk Leasequery operation (RFC 5460 [62]). Specifically, the CMTS SHOULD support query-by-remote-ID query type to query the DHCP server regarding leases assigned to devices behind a specific CM identified by its 48-bit MAC address.

10.2.5.3 Alternate Provisioning Mode (APM) Operation

When provisioning in Alternate Provisioning Mode, the CM tries to provision using IPv6 first. If IPv6 provisioning is unsuccessful, either because IPv6 Address acquisition or the TFTP configuration file download fails, the CM abandons IPv6 provisioning and attempts provisioning using IPv4. Figure 10-19 shows the process flow for APM.

If the CMTS has directed the CM to use APM and the IPv6 provisioning process fails, the CM MUST stop the IPv6 provisioning process.

If the CMTS has directed the CM to use APM and the IPv6 provisioning process fails, the CM MUST discard any provisioning information obtained up to that point in the provisioning process.

If the CMTS has directed the CM to use APM and the IPv6 provisioning process fails, the CM MUST release any IP addresses assigned up to that point in the provisioning process.

If the CMTS has directed the CM to use APM and the IPv6 provisioning process fails, the CM MUST note the event that IPv6 provisioning has failed in the Local Event Log.

If the CMTS has directed the CM to use APM and the IPv6 provisioning process fails, the CM MUST restart IP provisioning with the IPv4 provisioning mechanism described in clause 10.2.5.1. If the subsequent IPv4 provisioning fails, the CM MUST note the event that IPv4 provisioning has failed in the Local Event Log and scan for another downstream channel.

10.2.5.4 Dual-stack Provisioning Mode (DPM)

In Dual-stack Provisioning Mode (DPM), the CM attempts to acquire both IPv6 and IPv4 addresses and parameters through DHCPv6 and DHCPv4 almost simultaneously. For the acquisition of time-of-day and the download of a configuration file the CM prioritizes the use of the IPv6 address over the IPv4 address. If the CM cannot obtain an IPv6 address or if it cannot download a configuration file using IPv6, it tries downloading it using IPv4. In this mode, the CM makes both the IPv4 and the IPv6 addresses, if successfully acquired, available for management. Figure 10-20 shows the process flow for DPM.

When the CM is configured for DPM, its DHCPv4 and DHCPv6 clients operate independently. For example, the lease times for the IPv4 and IPv6 addresses may be different and the DHCP clients need not attempt to extend the leases on the IP addresses simultaneously.

If the CM is directed through the MDD message to operate in Dual-stack mode, the CM MUST perform IPv6 network connectivity as specified in clause 10.2.5.2. The CM MUST also perform IPv4 network connectivity as specified in clause 10.2.5.1. The CM MAY perform IPv4 network connectivity in parallel or after it has successfully obtained an IPv6 address. However, the CM MUST initiate the establishment of IPv4 network connectivity before attempting to acquire the current time of day with ToD over IPv6.

The CM MUST attempt to download a configuration file with IPv6 first. If the CM fails to acquire an IPv6 address, the CM MUST use TFTP over IPv4 for the download of a configuration file and log the event. If after acquiring an IPv6 address the CM fails to download a configuration file with TFTP over IPv6, the CM MUST log the event and attempt downloading a configuration file using TFTP over IPv4. If this attempt fails, the CM MUST log the event and scan for another downstream channel.

10.2.5.5 Establish Time of Day

The CM acquires time of day for the purpose of time stamping warning and error logs and messages and may also acquire it for the correct operation of some eSAFE devices. The CM acquisition of time is not required for successful CM provisioning.

The CM MUST attempt to obtain the current date and time by using the Time Protocol [27], as shown in figure 10-21. If the *Time Server Option* field is missing or invalid, the CM MUST initialize the current time to Jan 1, 1970, 0h00. In this case the CM MUST ignore the value, if any, of the Time Offset option.

The CM MUST use its DHCP-provided IP address for exchange of messages with the Time Protocol server. The CM MUST transmit the request using UDP [26]. The CM MUST listen for the response on the same UDP port as is used to transmit the request. The CM MUST combine the time retrieved from the server (which is UTC) with the time offset received from the DHCP server to create a notional "local" time.

The DHCP server may return multiple IP addresses of Time Protocol servers. The CM MUST attempt to obtain time of day from all the servers listed until it receives a valid response from any of the servers. Once the CM acquires time, it MUST stop requesting. The CM MUST contact the servers in batches of tries with each batch consisting of one try per server and each successive try within a batch at most one second later than the previous try and in the order listed by the DHCP message. If the CM fails to acquire time after any batch of tries, it MUST retry a similar batch using a truncated randomized binary exponential backoff with an initial backoff of 1 second and a maximum backoff of 256 seconds.

If a CM is unable to establish time of day before registration it MUST log the failure in the local log and, if configured for it, to syslog and SNMP trap servers. If the CM does not obtain ToD in the initial request against the first server, the CM MUST initialize the current time to Jan 1, 1970, 0h00 and then subsequently initialize its current time once it receives a response from a Time Server.

10.2.5.6 Transfer Operational Parameters

After the CM has attempted to obtain the time of day, the CM MUST download a configuration file using Trivial File Transfer Protocol (TFTP) ([31]), as shown in figure 10-21.

When using DHCPv4, if there are one or more addresses in the DHCPv4 TFTP Servers option in the DHCPACK, the CM MUST utilize the addresses listed in this option sequentially to obtain a configuration file and ignore the siaddr field. If there are no addresses provided in the DHCPv4 TFTP Servers option in the DHCPACK, the CM MUST use the address in siaddr to obtain a configuration file. The CM MUST use the name in the file field of the DHCPACK message to identify the configuration file to be downloaded.

When using DHCPv6, the CM MUST sequentially utilize the list of addresses in the TFTP Server Addresses option in the DHCPv6 Reply messages and MUST use the name in the Configuration File Name option in the Vendor Specific Information Options in the DHCPv6 Reply messages to identify the configuration file to be downloaded.

The CM follows the guidelines below in order to obtain a configuration file from the TFTP server:

- The CM MUST include the TFTP Blocksize option [40] when requesting the configuration file.
- The CM MUST request a blocksize of 1 448 if using TFTP over IPv4. REQ4331 The CM MUST request a blocksize of 1 428 if using TFTP over IPv6.
- The CM MUST initiate a configuration file download by sending a TFTP Read Request message for the configuration file using the TFTP Server address(es) obtained in the TFTP Servers Option or SIADDR, thus establishing a connection with the server [31]. When multiple TFTP Servers are present in the TFTP Servers Option, the CM progresses sequentially through the list of server IP addresses, attempting to successfully download a configuration file from each IP address until all retries and backoffs are exhausted for each of the server IP addresses.
- If the CM receives no response to the TFTP Read Request message, the CM MUST resend the TFTP Read Request up to TFTP Request Retries limit as defined in annex B.
- The CM MUST use an adaptive timeout between retries based on a binary exponential backoff with an initial backoff value of TFTP Backoff Start and final backoff value of TFTP Backoff End as defined in annex B:
 - The CM MUST log an event in the local log for each failed attempt.
 - If the CM receives no response to the TFTP Read Request after all of the TFTP Request Retries (annex B), the CM MUST restart the configuration file download process on the next server in the list of servers.
- The CM MUST attempt to download a configuration file from the first entry in the DHCPv4 TFTP Servers option list and exhaust all backoffs and retries before moving to the next entry in the list until successful reception of a configuration file.

The CM follows these general guidelines when provisioned for IPv6 operation:

- If the CM reaches the end of the TFTP Server Addresses option list before a successful download of a configuration file, the CM will declare IP Connectivity has failed.
- If the CM cannot download a valid configuration file from a TFTP server, either because the CM receives a TFTP error message from the TFTP server or because the configuration file downloaded is invalid as defined in clause 10.2.5.7, the CM MUST retry the configuration file download process up to the TFTP Download Retries (annex B) after waiting TFTP Wait time (annex B) without performing the TFTP Read Request Retries (annex B).
- If the CM cannot download a valid configuration file, as in the previous bullet, after all of the TFTP Download Retries (annex B), the CM MUST restart the configuration file download process on the next server in the list of servers.

If the CM receives an ICMP Destination Unreachable message for the current TFTP server at any time during the configuration file download process, the CM MUST terminate the configuration file download on the TFTP server whose address is included in the ICMP Destination Unreachable message without performing the TFTP Read Request Retries or the TFTP Download Retries (annex B). The CM MUST restart the configuration file download process on the next server in the list of servers.

If the CM reaches the end of the TFTP Server Addresses option list before a successful download of a configuration file, the CM MUST declare IP Connectivity has failed and log an event.

10.2.5.7 Configuration File Processing

After downloading the configuration file and prior to completing IP Provisioning, the CM performs several processing steps with the configuration file.

If a modem downloads a configuration file containing an Upstream Channel ID Configuration Setting (clause C.1.1.2) different from what the modem is currently using, the modem MUST NOT send a Registration Request message to the CMTS. Likewise, if a modem downloads a configuration file containing a Single Downstream Channel Frequency (clause C.1.1.22.1.2) and/or Downstream Frequency Range (clause C.1.1.22.2) that does not include the downstream frequency the modem is currently using or a Downstream Frequency Configuration Setting (clause C.1.1.1) different from what the modem is currently using and no Downstream Channel List, the modem MUST NOT send a Registration Request message to the CMTS. In either case, the modem MUST redo initial ranging using the configured upstream channel and/or downstream frequency(s) per clause 10.2.3.

If the modem downloads a configuration file containing the Enable 2.0 Mode TLV set to Disable (see clause C.1.1.19.5):

- If the CM is using a Type 3 or a Type 4 channel, it MUST NOT send a Registration Request message to the CMTS. In this case, the CM MUST redo initial ranging using a Type 1 or Type 2 channel. If no such upstream channel is available (or if the modem is unable to range successfully on one) the modem MUST scan for a new downstream (at which point the CM may once again operate in 2.0 Mode).
- If the CM is using any other type channel, the CM MUST continue with the IP Provisioning and Registration processes.
- If Multiple Transmit Channel Mode is disabled during registration, then the CM MUST NOT operate in 2.0 Mode until it registers again and downloads a configuration file which does not have this TLV set to Disable.
- If Multiple Transmit Channel Mode is enabled during registration, then the CM MUST ignore the value of the Enable 2.0 Mode TLV and use any Type of upstream assigned by the CMTS that is supported by the modem.

If the modem downloads a configuration file containing the Enable 2.0 Mode TLV set to Enable (see clause C.1.1.19.5) or that does not contain the Enable 2.0 Mode TLV:

- If Multiple Transmit Channel Mode is disabled during registration, then the CM MUST enable 2.0 Mode until it registers again and downloads a configuration file which does not have this TLV set to Disable. This means that, if a modem registers on a Type 1 channel with a configuration file that enables 2.0 Mode operation and then ends up on a Type 2, Type 3 or Type 4 upstream without going through re-registration, the CM would immediately start operating in 2.0 Mode.
- If Multiple Transmit Channel Mode is enabled during registration, then the CM MUST ignore the value of the Enable 2.0 Mode TLV and use any Type of upstream assigned by the CMTS that is supported by the modem.

The CM performs additional operations to verify the validity of a configuration file and MUST reject a configuration file that is invalid. An invalid configuration file is a file with any of these characteristics:

- Lacks one or more mandatory items, as defined in clause D.1.2.
- Has an invalid MIC, as defined in clause D.1.3.1.
- Has one or more TLV-11 encodings that cannot be processed and cause rejection of the file, as defined in [10].
- Contains a TLV-53 encoding, SNMPv1v2c Coexistence Configuration, that causes rejection of the file, as defined in clause C.1.2.13.
- Contains a TLV-54 encoding, SNMPv3 Access View Configuration, that causes rejection of the file, as defined in clause C.1.2.14.
- Contains a TLV-60 encoding, Upstream Drop Classifiers, that has an invalid value or length.

The CM MUST NOT reject a configuration file unless it is considered as invalid under conditions specified above. The CM MUST continue with Registration Request under conditions other than specified above.

10.2.5.8 Post-registration Failures to Renew IP Addresses

If the CM is configured to provision in IPv4 Only or IPv6 Only mode and it fails to renew its IP address it MUST reinitialize the MAC as defined in clause 10.2.

If a CM is configured to use APM and the CM fails to renew its IP address, the CM MUST note the event in the Local Event Log. Failure to renew the IP address is a critical event. After noting the failure in the Local Event Log, the CM MUST reinitialize the MAC as defined in clause 10.2.

If a CM is configured to use DPM and the CM fails to renew one of its IP addresses, the CM MUST note the event in the Local Event Log. Failure to renew an IP address when the other IP address is active is not a critical event. In this case, after noting the failure in the Local Event Log, the CM MUST continue to operate with the remaining IP address. On the other hand, failure to renew an IP address is a critical event when the other IP address has already expired. When a CM operating in DPM fails to successfully renew its only remaining IP address, the CM MUST reinitialize the MAC as defined in clause 10.2.

10.2.6 Registration with the CMTS

10.2.6.1 Cable Modem Requirements

Once the CM establishes IP Connectivity and unless directed to a different Primary Downstream Channel via the configuration file (see clauses C.1.1.1 and C.1.1.22), the CM MUST register with the CMTS per figure 10-25 through figure 10-30. In this clause, the term Registration Request refers to either the REG-REQ MAC Management Message or the REG-REQ-MP MAC Management Message. The term Registration Response refers to either the REG-RSP MAC Management Message or the REG-RSP-MP MAC Management Message.

- 1) The CM creates a Registration Request message which includes all its CM capabilities and the CM Receive Channel Profile(s). If the Primary DS Channel contained an MDD message, the CM creates a REG-REQ-MP. Otherwise the CM creates a REG-REQ message. The CM sends the REG-REQ or REG-REQ-MP message to the CMTS, starts timer T6 and then awaits a response.
- 2) If the CM receives a fragment of a REG-RSP-MP message, the CM returns to the Waiting for REG-RSP or REG-RSP-MP state and waits for the next fragment. Once the CM has received a REG-RSP or all the REG-RSP-MP fragments, it stops the T6 timer. If the T6 timer expires before the REG-RSP or all fragments of the REG-RSP-MP is received, the CM retransmits the REG-REQ or REG-REQ-MP. Upon reaching the retransmission limit annex B, the CM will perform a MAC reinitialization.
- 3) If the CM receives a REG-RSP or all the REG-RSP-MP fragments before timer T6 expires and the response is not equal to "okay", the CM will either send a REG-ACK with the appropriate error sets and then perform a MAC reinitialization or it will perform the MAC reinitialization directly. The CM MUST send the REG-ACK if:
 - a) the CM sent a REG-REQ-MP;
 - b) the REG-RSP contained QoS (Service Flow) encodings;
 - c) the CM is operating on a Type 3 or Type 4 upstream channel; or
 - d) the CM is operating on a Type 2 channel and "DOCSIS 2.0 Mode" is not disabled in the CM config file.

If none of those conditions are true, then the CM will not send the REG-ACK. In DOCSIS 1.0, the CM was not required to send a REG-ACK after receiving a REG-RSP.

- 4) The CM checks the Registration Response and verifies that all parameters can be supported. After processing the Registration Response, the CM MUST NOT transmit upstream traffic until it sends the REG-ACK. If one or more parameters cannot be supported, the CM sends a REG-ACK with the appropriate error sets of the unsupported parameters.

As a part of verifying all parameters, the CM checks that the RCC and TCC encodings (if present) are consistent with the CM's hardware capabilities. If the RCC or TCC encodings are not consistent, the CM sends a REG-ACK with an error code of "reject-bad-rcc" or "reject-bad-tcc" encodings (refer to clause C.4). The CM will then perform a MAC reinitialization with a CM Initialization Reason of BAD_RCC_TCC after sending the REG-ACK message.

- 5) If RCC encodings are not present in the Registration Response, the CM transitions to the CM Complete Registration state. If RCC encodings are present in the Registration Response, then the CM will attempt to acquire all the receive channels in the RCC. If the CM supports multiple receive channels, the CM transitions to the AcquireDS(RC) subroutine for each downstream receive channel. The CM attempts FEC, MPEG and SYNC lock on the Primary Downstream Channel and FEC and MPEG lock on the non-primary downstream receive channels.

If the downstream acquisition fails on the primary downstream, the CM aborts all other receive channel acquisition processes and saves the "Failed Primary DS" state information. The CM then performs a MAC reinitialization with a CM Initialization Reason of FAILED_PRIM_DS.

If the downstream acquisition was successful on the primary downstream but failed on one of the other downstream channels in the RCC encodings, the CM begins operating in a partial service mode of operation in the downstream, sets the REG-ACK error code to "partial-service" (refer to clause C.4) and proceeds to acquire the transmit channels.

If the downstream acquisition was successful on all the downstream channels in the RCC encodings, the CM proceeds to acquire the transmit channels.

- 6) If TCC encodings are not present in the Registration Response, the CM transitions to the CM Complete Registration state.

If TCC encodings are present in the Registration Response, the CM transitions to the AcquireUS(TC) subroutine. If the TCC Upstream Channel Action (refer to clause C.1.5.1.2) is equal to "no action", the upstream channel is the channel on which the Registration Request message was sent and the CM continues to range with the Temporary SID becoming the Ranging SID. Otherwise, the CM attempts ranging on all the upstream channels, per the Ranging SID and TCC initialization technique encodings. If the CMTS does not explicitly include the channel on which the Registration Request message was sent in the TCC Encodings with a TCC Upstream Channel Action of "no action", "change", or "delete" or implicitly include the channel on which the Registration Request message was sent in the TCC Encodings with a TCC Upstream Channel Action of "re-range", the CM considers the Registration Response to be invalid. If the CMTS does not include a TCC encoding with an Upstream Channel Action of Re-range in the Registration Response when the Receive Channel Center Frequency Assignment (subtype 49.5.4) of the Primary Downstream is not the same as the Receive Module First Channel Center Frequency (subtype 49.4.4) of the Receive Module containing the Primary Downstream and one of the upstream channels assigned in the TCC encoding is an S-CDMA channel, the CM rejects the Registration Response and performs a MAC reinitialization with a CM Initialization Reason of BAD_RCC_TCC.

If the "Acquire CM Transmit Channels" subroutine fails to range on all the upstream channels in the TCS, the CM needs to save the "TCS Failed on All Upstream Channels" state information. The CM then performs a MAC reinitialization with a CM Initialization Reason of TCS_FAILED_ON_ALL_US.

If Multiple Transmit Channel Support is enabled and the CM is able to successfully range on one or more (but not all) of the upstream channels in the TCS, the CM will start Multiple Transmit Channel Mode (refer to clause C.1.3.1.24) in a partial service mode of operation in the upstream. In this case, the CM will set the REG-ACK error code to "partial-service". (refer to clause C.4).

If Multiple Transmit Channel Support is enabled and the CM is able to successfully range on all of the upstream channels in the TCS, the CM will start Multiple Transmit Channel Mode (refer to clause C.4). In this case, the CM will set the REG-ACK confirmation code set to "okay".

If Multiple Transmit Channel Support is zero (disabled) and the CM was able to successfully range on its one upstream channel, the CM does not enable Multiple Transmit Channel Mode.

If the "Acquire CM Transmit Channels" subroutine returns a TCS Success or a TCS Partial Service, the CM proceeds to the "CM Complete Registration" process.

- 7) In the CM Complete Registration state, the CM sets up the service flows, assigns SID Clusters for available transmit channels and activates all operational parameters.

The CM can create the primary upstream service flow if and only if it successfully ranges on at least one of the upstream channels defined in the SID-to-Channel Mapping SID Cluster encoding.

If the CM can create the primary upstream service flow, the CM sends the REG-ACK message to the CMTS and starts the T10 transaction timer.

If the CM cannot create the primary upstream service flow, then the CM will not send a REG-ACK and will perform a MAC reinitialization with a CM Initialization Reason of NO_PRIM_SF_USCHAN.

- 8) If Multiple Transmit Channel Mode is disabled and if DOCSIS 2.0 mode is enabled (refer to clause C.1.1.19.5) and if the CM is transmitting on a Type 2 or Type 4 upstream channel (refer to clause 6.4.3), then the CM transitions to the REG-HOLD2 state. If the T10 transaction timer expires, the CM begins operating in DOCSIS 2.0 mode and uses IUCs 9, 10 and 11 for data transmission. Note that this operation differs slightly from DOCSIS 2.0 initialization.

If the CM receives a Registration Response in the REG-HOLD2 state (e.g. due to the CMTS not receiving the REG-ACK), the CM sends another REG-ACK and restarts the T10 timer.

- 9) If Multiple Transmit Channel Mode is enabled or, if Multiple Transmit Channel Mode is disabled and DOCSIS 2.0 mode is disabled (refer to clause C.1.1.19.5) or, if Multiple Transmit Channel Mode is disabled and DOCSIS 2.0 mode is enabled and the CM is not on a Type 2 or Type 4 upstream, then the CM completes registration and transitions to the REG-HOLD1 state. If the CM receives a Registration Response message while in the REG-HOLD1 state prior to the expiration of the T18 timer, (e.g. due to the CMTS not receiving the REG-ACK), the CM retransmits the REG-ACK starts the T10 timer and re-enters the REG-HOLD1 state. If the CM then receives another Registration Response message while in the REG-HOLD1 state prior to the expiration of the T10 Timer, (e.g. due to the CMTS not receiving the REG-ACK), the CM retransmits the REG-ACK, re-starts the T10 timer and re-enters the REG-HOLD1 state.

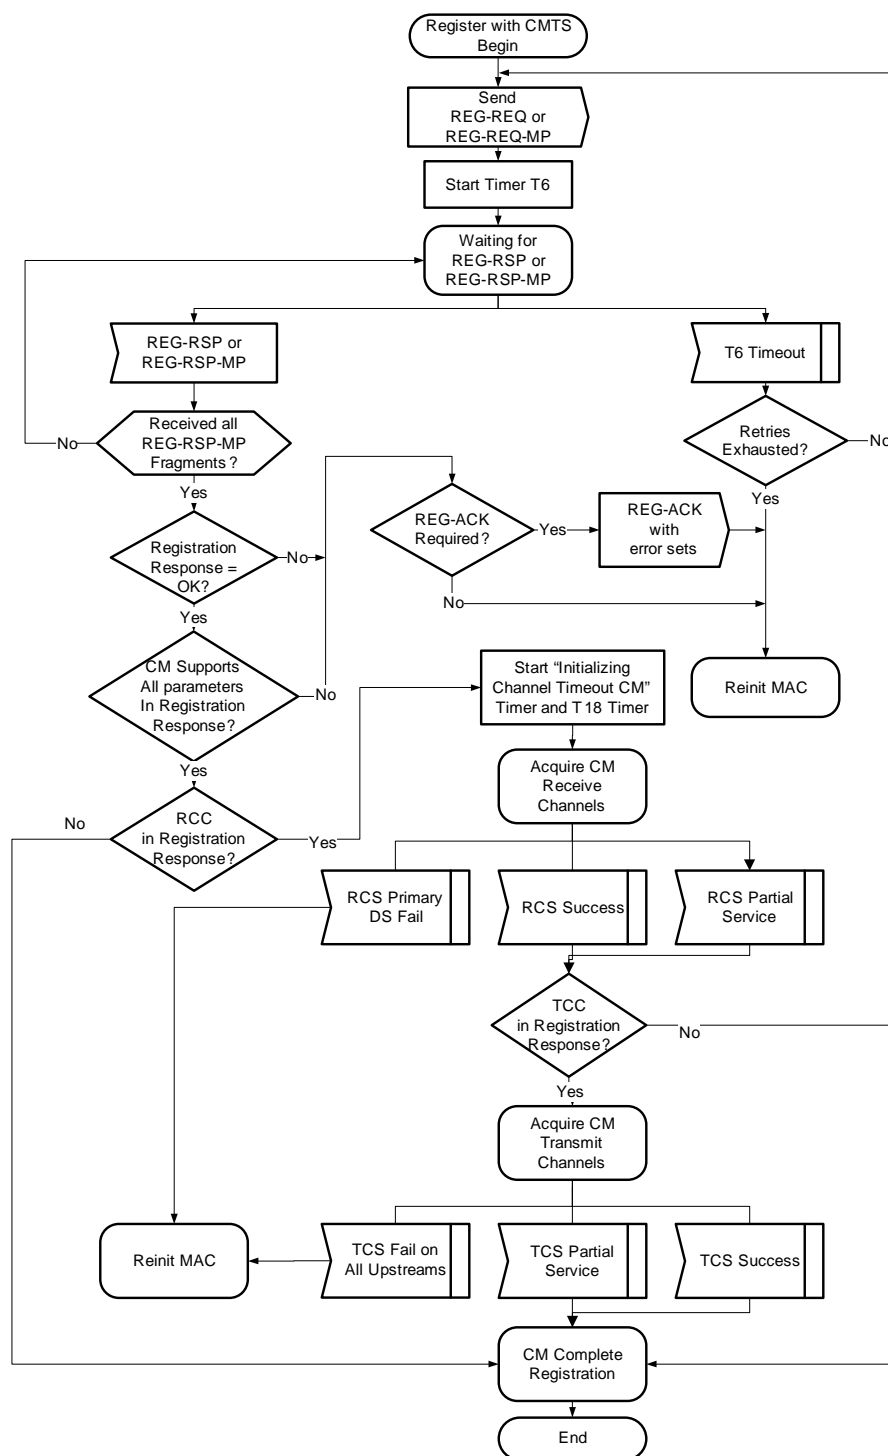


Figure 10-25: CM Register with CMTS - Begin

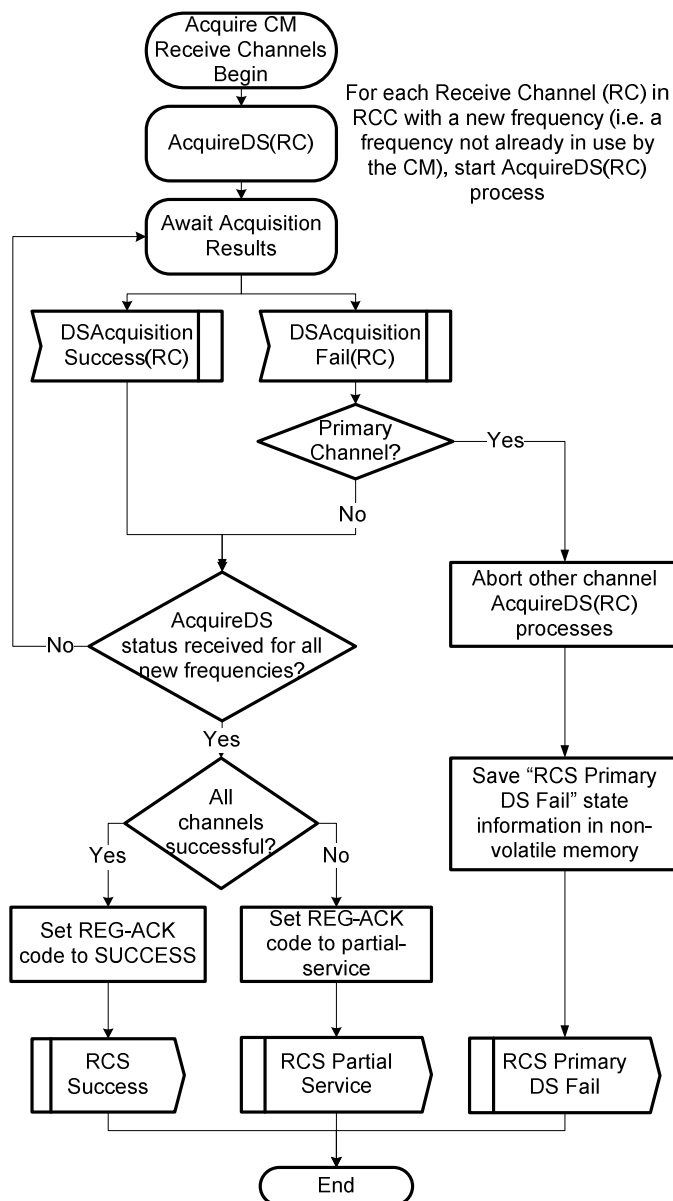


Figure 10-26: CM Acquires Receive Channels

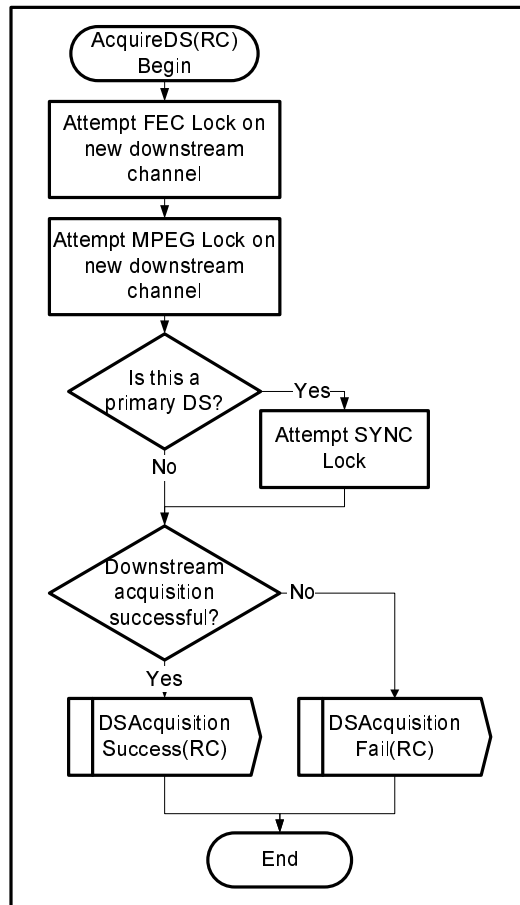


Figure 10-27: CM Acquires Downstream Channels

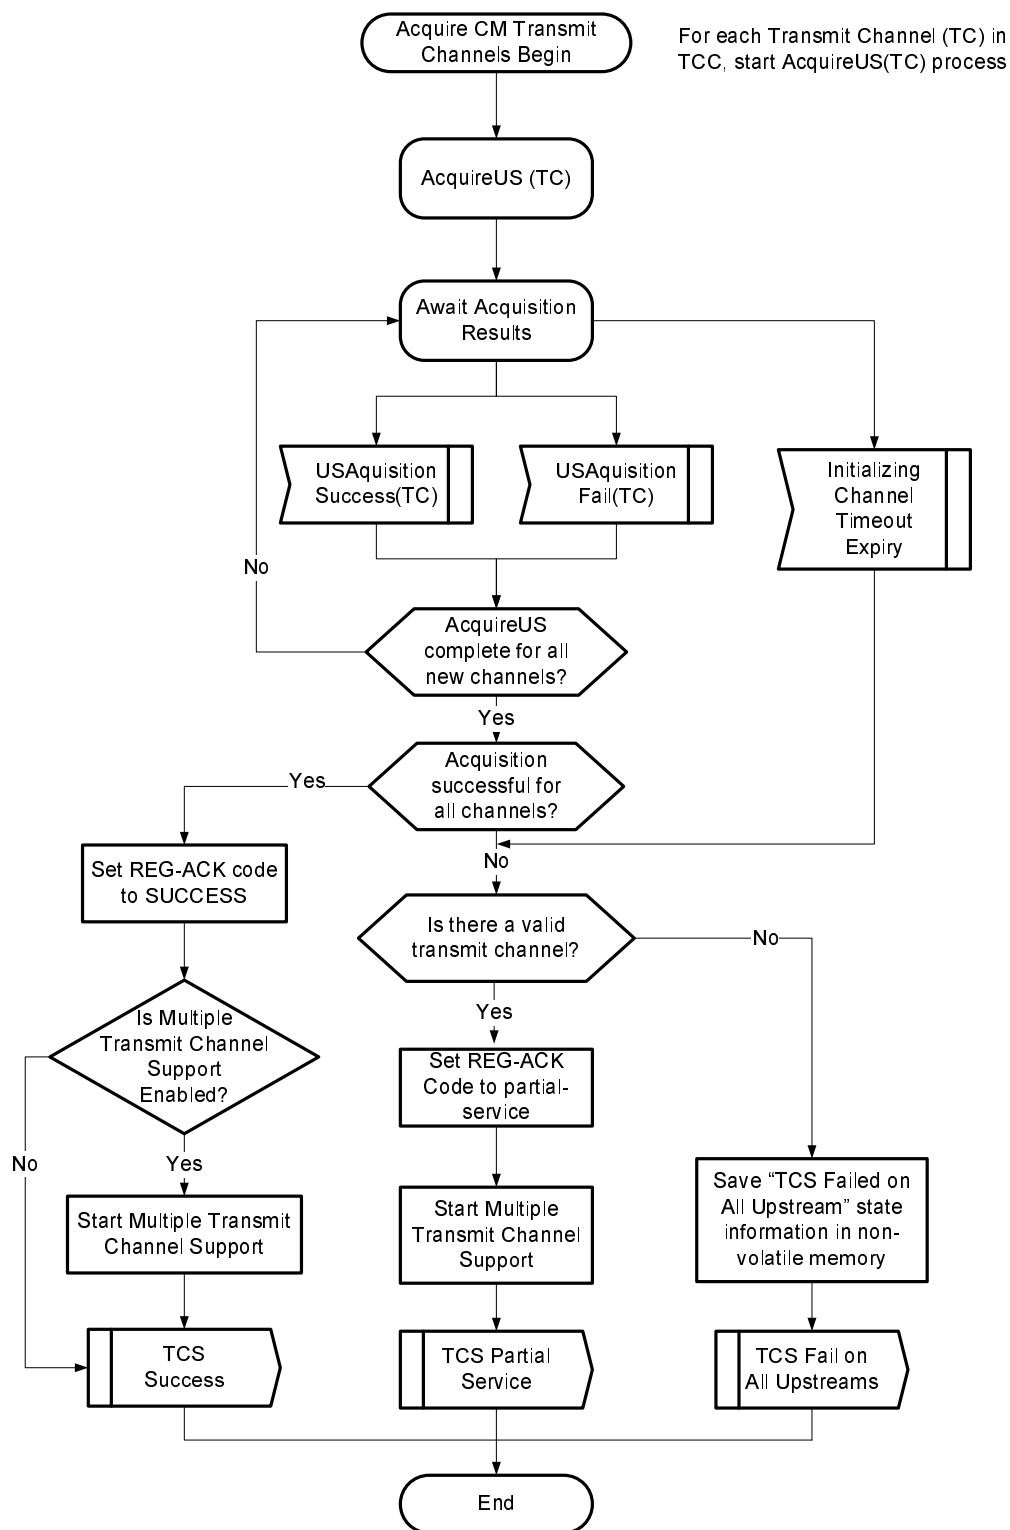
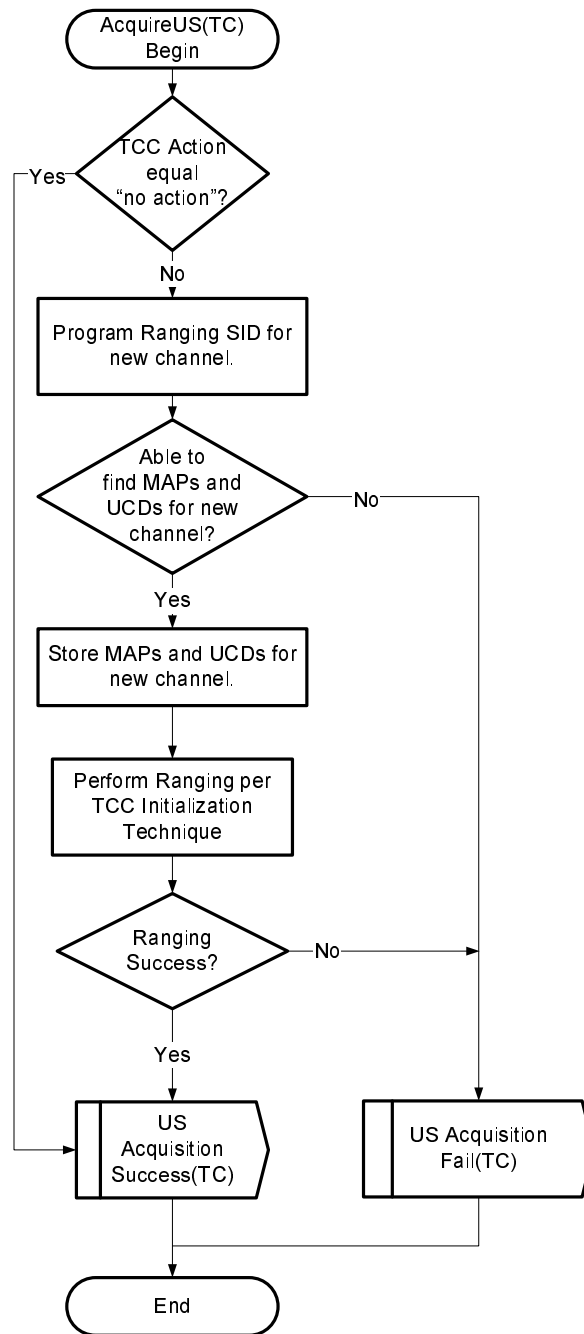


Figure 10-28: CM Acquires Transmit Channels

**Figure 10-29: CM Acquires Upstream Channels**

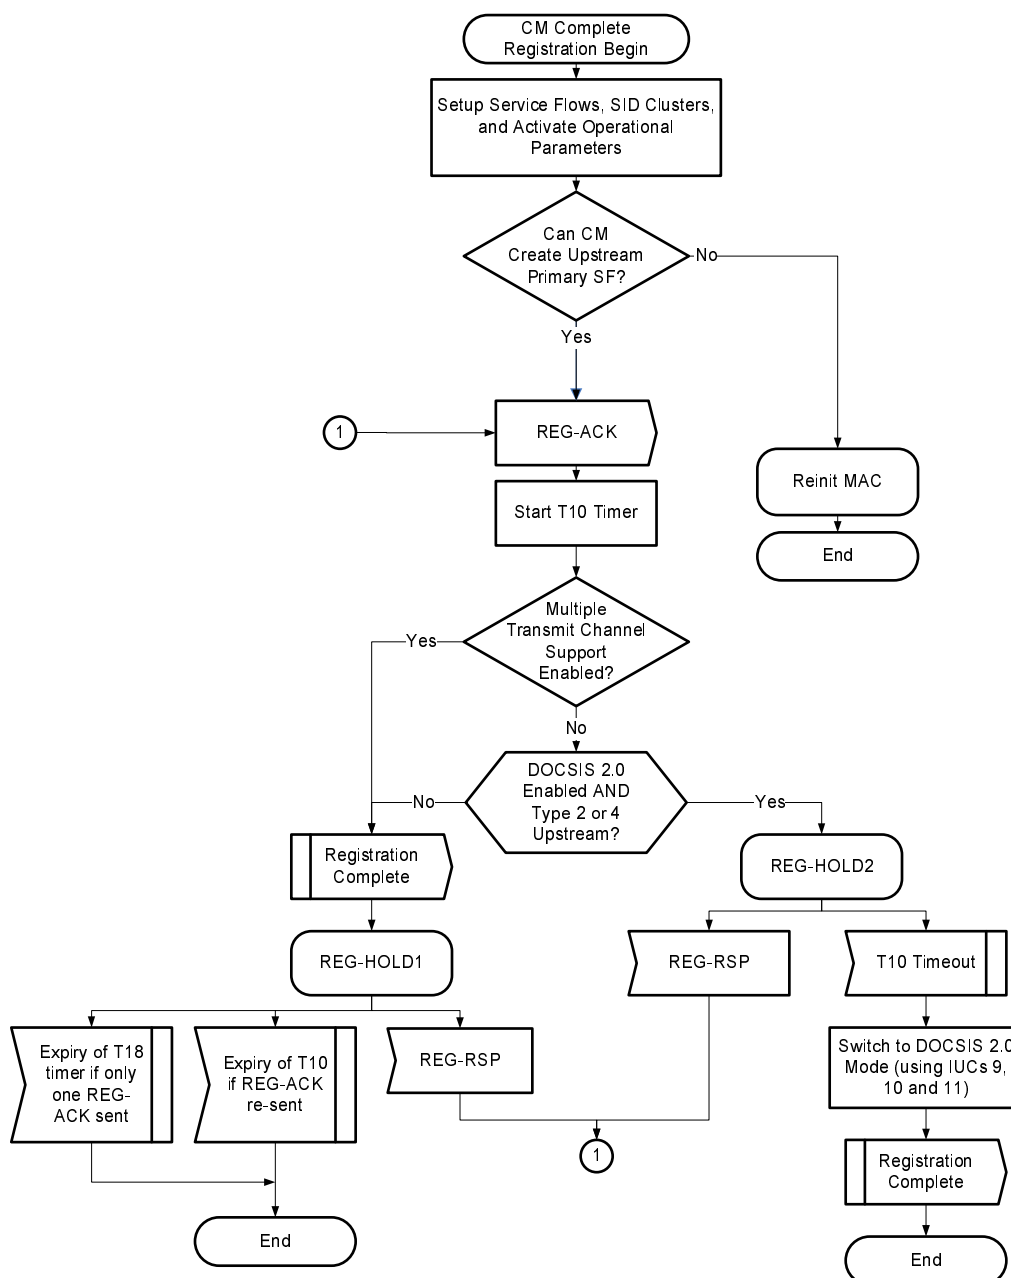


Figure 10-30: CM Completes Registration

10.2.6.2 CMTS Requirements

In this clause, the term Registration Request refers to either the REG-REQ MAC Management Message or the REG-REQ-MP MAC Management Message. The term Registration Response refers to either the REG-RSP MAC Management Message or the REG-RSP-MP MAC Management Message. If the CM sent a REG-REQ message, the CMTS responds with a REG-RSP message. If the CM sent a REG-REQ-MP message, the CMTS responds with a REG-RSP-MP message.

The CMTS MUST perform the CM registration process as shown in figure 10-31 though figure 10-33.

- 1) The CMTS waits for the REG-REQ or REG-REQ-MP message. If timer T9 expires, the CMTS de-assigns the temporary SID for the CM and makes some provision for aging out the temporary SID.

- 2) If the CMTS receives a fragment of a REG-REQ-MP message, the CMTS returns to the Waiting for REG-REQ or REG-REQ-MP state and waits for the next fragment. Once the CMTS has received a REG-REQ or all REG-REQ-MP message fragments, it stops the T9 timer and proceeds with MIC calculations. Note that the CMTS is required to send multipart registration response messages in response to receiving multipart registration request messages (refer to clause 6.4.8.2).
- 3) The CMTS performs the MIC procedures defined in clause D.2.1. If MIC verification fails, the CMTS responds with an Authentication Failure in the REG-RSP.
- 4) If the TFTP Server Timestamp field is present, the CMTS checks if the time is different from its local time by more than CM Configuration Processing Time (refer to annex B). If the time is different, then the CMTS MUST indicate Authentication Failure in the Response field of the Registration Response. The CMTS SHOULD also log an entry listing the CM MAC address from the message.
- 5) If the TFTP Server Provisioned Modem Address field is present, the CMTS checks if the Provisioned Modem Address matches the requesting modem's actual address. If the addresses do not match, the CMTS MUST indicate Authentication Failure in the Response field of the Registration Response. The CMTS SHOULD also log an entry listing the CM MAC address from the message.
- 6) If the CMTS cannot support the requested services, it sends the Registration Response with the appropriate non-zero confirmation code (refer to clause C.4) and terminates processing of the Registration Request.

If the Registration Request contained DOCSIS 1.0 Class of Service encodings, the CMTS verifies the availability of the class(es) of service requested. If the CMTS is 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 clause C.1.1.4).

If the Registration Request contained Service Flow encodings, the CMTS verifies the availability of the Quality of Service requested in the provisioned Service Flow(s). If the CMTS is unable to provide the Service Flow(s), the CMTS MUST respond with either a reject-temporary or reject-permanent (refer to clause C.4) and the appropriate Service Flow response codes.

If the Registration Request contained both 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 Classes and Service Flows requested.

When the CMTS sends a REG-RSP with a non-zero confirmation code and the CM is not expected to send a REG-ACK, the CMTS will return to the Waiting for REG-REQ state. Otherwise, if the CMTS sends the REG-RSP or REG-RSP-MP with a non-zero confirmation code, the CMTS waits for the REG-ACK from the CM.

Note that the CM will send the REG-ACK if:

- a) the CM sent a REG-REQ-MP;
- b) the REG-RSP contained QoS (Service Flow) encodings;
- c) the CM is operating on a Type 3 or Type 4 upstream channel; or
- d) the CM is operating on a Type 2 channel and "DOCSIS 2.0 Mode" is not disabled in the CM config file.

If none of those conditions are true, then the CM will not send the REG-ACK.

- 7) The CMTS then verifies the availability of all Modem Capabilities requested. If unable or unwilling to provide the Modem Capability requested, the CMTS turns off (sets to 0) that Modem Capability (refer to clause 6.4.8.3.1) in the Registration Response.
- 8) If the CM supports multiple receive channels, the CMTS will check the Receive Channel Profile (RCP). If the CM indicated support for multiple downstream receive channels in a REG-REQ-MP that did not include an RCP, the CMTS returns a REG-RSP-MP with "Missing RCP error" (refer to clause C.4) and terminates processing of the REG-REQ-MP. The CMTS then waits for the REG-ACK from the CM.

If the CM indicated support for multiple receive channels in a REG-REQ (as opposed to a REG-REQ-MP), the CMTS SHOULD disable Multiple Receive Channel mode by returning a value of zero for Multiple Receive Channel Support in the REG-RSP.

- 9) If the CM supports multiple receive channels and included one or more RCP encodings in the REG-REQ-MP, the CMTS either enables Multiple Receive Channel Mode (by confirming the value in the REG-REQ-MP) or disables Multiple Receive Channel Mode by setting the capability to zero in the REG-RSP-MP.

If the CMTS enables Multiple Receive Channel Mode, the CMTS MUST populate a Receive Channel Configuration (RCC) Encoding in the REG-RSP. The RCC encoding configures the CM's physical layer components to specific downstream frequencies.

If the Receive Channel Center Frequency Assignment (subtype 49.5.4) of the Primary Downstream is not the same as the Receive Module First Channel Center Frequency (subtype 49.4.4) of the Receive Module containing the Primary Downstream and the TCC encoding contains at least one SCDMA channel, the CMTS includes a TCC encoding with an Upstream Channel Action of Re-range in the Registration Response.

- 10) If the CM included a non-zero Multiple Transmit Channel Support capability in the REG-REQ-MP, the CMTS either enables Multiple Transmit Channel Mode (by confirming the value in the REG-REQ-MP) or disables Multiple Transmit Channel Mode by setting the capability to zero in the REG-RSP-MP. If the CM does not support Multiple Receive Channel mode or the CMTS disabled Multiple Receive Channel mode in step 8 or 9, the CMTS is required to disable Multiple Transmit Channel mode (clause C.1.3.1.29). If the CM included a Multiple Transmit Channel Support TLV with a value of 0, the CMTS returns a value of 0 for Multiple Transmit Channel Support in the Registration Response, disabling Multiple Transmit Channel Mode.

If the CMTS enables Multiple Transmit Channel Mode, the CMTS MUST include TCC encodings in the REG-RSP-MP. If the CMTS enables Multiple Receive Channel Mode and disables Multiple Transmit Channel Mode (by setting or confirming a capability of zero), the CMTS MAY include TCC encodings for the single upstream channel in the REG-RSP-MP. A TCC encoding defines the CM operations to be performed on an upstream channel in the Transmit Channel Set. When the CMTS sends a TCC encoding in the REG-RSP-MP, the CMTS MUST subsequently use DBC signalling (as opposed to DCC or UCC messaging) to make changes to the TCS. When the CMTS does not assign a Transmit Channel Configuration during registration, the CMTS may use DCC signalling to change to the upstream channel.

If the CMTS includes TCC encodings in the REG-RSP-MP and load balancing is disabled for the CM, the CMTS cannot delete or replace the CM's current upstream change (see clause 11.6.4).

When the CMTS includes TCC encodings in the REG-RSP-MP, the CMTS MUST also include Service Flow SID Cluster Assignments in the REG-RSP-MP. When Service Flow SID Cluster assignments are included in the REG-RSP-MP, the CMTS MUST NOT include a SID assignment under the Service Flow encodings portion of the REG-RSP-MP. If Multiple Transmit Channel Mode is disabled and TCC/SF SID Cluster Assignment encodings are included in the REG-RSP-MP, the CMTS MUST include only a single Service Flow SID Cluster assignment corresponding to the single channel in the TCC. In this case, the CMTS MUST set the Ranging SID to be the same as the SID corresponding to the Primary Upstream Service Flow.

When the CMTS includes TCC encodings in the REG-RSP-MP, the CMTS MUST include the upstream channel on which the CM transmitted the Registration Request message in the TCC encodings explicitly with a TCC Upstream Channel Action of "no action", "change", or "delete" or implicitly with an Upstream Channel Action of "re-range". If the CM is to continue transmitting on the upstream channel on which it transmitted the Registration Request message and to continue ranging with the temporary SID becoming the Ranging SID, the CMTS includes the upstream channel in the TCC encodings with a TCC Upstream Channel Action (refer to clause C.1.5.1.2) equal to "no action". If the CM is to continue transmitting on the upstream channel on which it transmitted the Registration Request message using a new Ranging SID, the CMTS includes the upstream channel in the TCC encodings with a TCC Upstream Channel Action equal to "change". If the upstream channel on which the CM transmitted the Registration Request message is not to be a part of the TCC, the CMTS includes this channel in the TCC encodings and uses a TCC Upstream Channel Action of "delete".

If the CMTS makes a change to the CM's Primary Downstream Channel, the CMTS MUST include a TCC encoding with an Upstream Channel Action of Re-range in the Registration Response. If the CMTS makes a change which affects the CM's Primary Downstream Channel and the TCC encoding contains at least one S-CDMA channel, the CMTS MUST include a TCC encoding with an Upstream Channel Action of Re-range in the Registration Response. This means that the CMTS cannot change the Primary Downstream Channel in the RCC during registration unless a TCC encoding has also been included in the REG-RSP-MP.

If the CM did not send a Multiple Transmit Channel Support capability in the Registration Request or the CMTS did not enable Multiple Receive Channel mode, the CMTS MUST NOT send TCC encodings in the Registration Response. In this case, the CMTS MUST NOT use DBC signalling to change upstream channels.

- 11) If the CMTS has enabled Multicast DSID Forwarding, on the CM, the CMTS assigns the appropriate DSIDs and Security Associations (SAIDs) to the multicast flows (refer to clause 9).
- 12) The CMTS creates all requested services, assigns Service Flows to channel sets and assigns Service Flow IDs, as appropriate. If the CMTS includes TCC Encodings in the Registration Response, the CMTS populates the Service Flow SID Cluster assignments in the REG-RSP-MP; if the "Initializing Channel Timeout" is different than the default value, the CMTS will populate the timer in the REG-RSP-MP.
- 13) If the CMTS includes a TCC in the REG-REQ-MP, the CMTS starts the "Initializing Channel Timeout" timer and sends the REG-RSP-MP with a confirmation code of okay(0). If no TCC is included, the CMTS starts the T6 timer and sends the Registration Response with a confirmation code of okay(0).

If the Registration Response was sent with a confirmation code of okay(0) and the CM is expected to send a REG-ACK, the CMTS waits for the REG-ACK. If the CM is not expected to send a REG-ACK, the CMTS does not wait for the REG-ACK.

Note that the CM will send a REG-ACK if:

- a) the CM sent a REG-REQ-MP;
- b) the REG-RSP contained QoS (Service Flow) encodings;
- c) the CM is operating on a Type 3 or Type 4 upstream channel; or
- d) the CM is operating on a Type 2 channel and "DOCSIS 2.0 Mode" is not disabled in the CM config file.

If none of those conditions are true, then the CM will not send the REG-ACK.

- 14) Once the CMTS sends the Registration Response, it waits for a bandwidth request from the CM. If the CMTS receives a non Queue-depth Based Request from the CM and Multiple Transmit Channel Mode is disabled, the CMTS continues to grant bandwidth to the CM using the non Queue-depth Based Request (legacy) mechanism.

If Multiple Transmit Channel Mode is enabled and the CMTS receives a non Queue-depth Based Request while waiting for the REG-ACK, the CMTS MUST support one of the following two options:

- 1) The CMTS ignores the request and waits for a Queue-depth Based Request from the CM. This might result in the CM re-initializing its MAC if the REG-RSP was lost.
 - 2) The CMTS grants the legacy request and goes back to waiting for a new REG-REQ or REG-REQ-MP. The CMTS is only allowed to accept legacy requests from a CM operating in Multiple Transmit Channel mode from the time the CMTS sends the REG-RSP until it receives a valid REG-ACK. After that, the CMTS ignores any non Queue-depth Based Requests received from that CM.
- 15) If the CMTS receives a Queue-depth Based Request and Multiple Transmit Channel Mode is enabled, the CMTS starts granting bandwidth to the CM using Multiple Transmit Channel Mode (refer to clause C.1.3.1.24). If the T6 timer expires while the CMTS is waiting for a REG-ACK after receiving the Queue-depth based request, the CMTS re-starts the T6 timer and re-sends the REG-RSP-MP message.

If the CMTS receives a Queue-depth Based Request and Multiple Transmit Channel Mode is disabled, the CMTS ignores the request and waits for a non Queue-depth Based Request from the CM.

NOTE 1: If the CMTS needs to change a CM's upstream operation to or from Multiple Transmit Channel Mode, the CMTS will force the CM to reinitialize and effect the change when the CM re-registers.

- 16) If the CMTS included a TCC in the REG-REQ-MP and the "Initializing Channel Timeout CMTS" timer expires, the CMTS clears any reassembly buffers, restarts the T6 timer and sends another Registration Response message. If no TCC was included, the T6 timer expires and all the Registration Response retries have not been exhausted, the CMTS clears any reassembly buffers, restarts the T6 timer and sends another Registration Response message. If the Registration Response retries have been exhausted, the CMTS destroys all services and Registration ends unsuccessfully.

NOTE 2: If the CMTS was using the "Initializing Channel Timeout CMTS" while waiting for the first REG-ACK, it still uses the T6 timer while waiting for a REG-ACK message from re-transmitted Registration Response messages.

- 17) Once the CMTS receives the REG-ACK message from the CM, it checks the message for error sets. If the REG-ACK includes only partial service as the error set, the CMTS may initiate action on the "partial service" using DBC signalling. The CMTS may try to reacquire failed channels, move the CM to other downstream or upstream channels or other CMTS specific alternatives. If the REG-ACK contains error sets other than partial service, the CMTS destroys all services and Registration ends unsuccessfully.
- 18) If the REG-ACK contains no error sets, DOCSIS 2.0 mode has been enabled (refer to clause C.1.1.19.5), Multiple Transmit Channel Mode is disabled and the upstream is a Type 2 or Type 4 channel, the CMTS MUST begin to use IUCs 9, 10 and 11 for all grants to the CM.

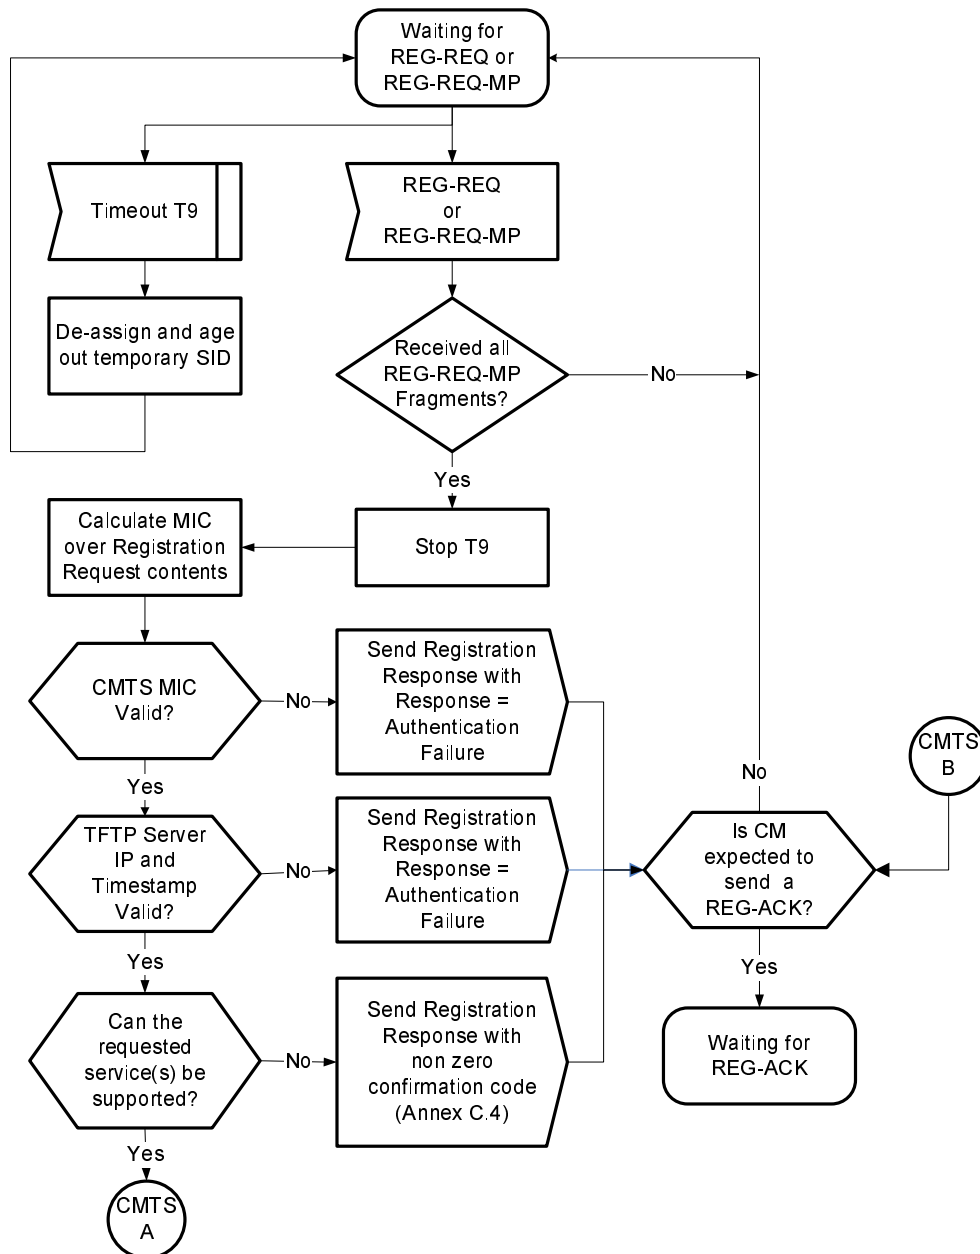


Figure 10-31: CMTS Registration - Begin

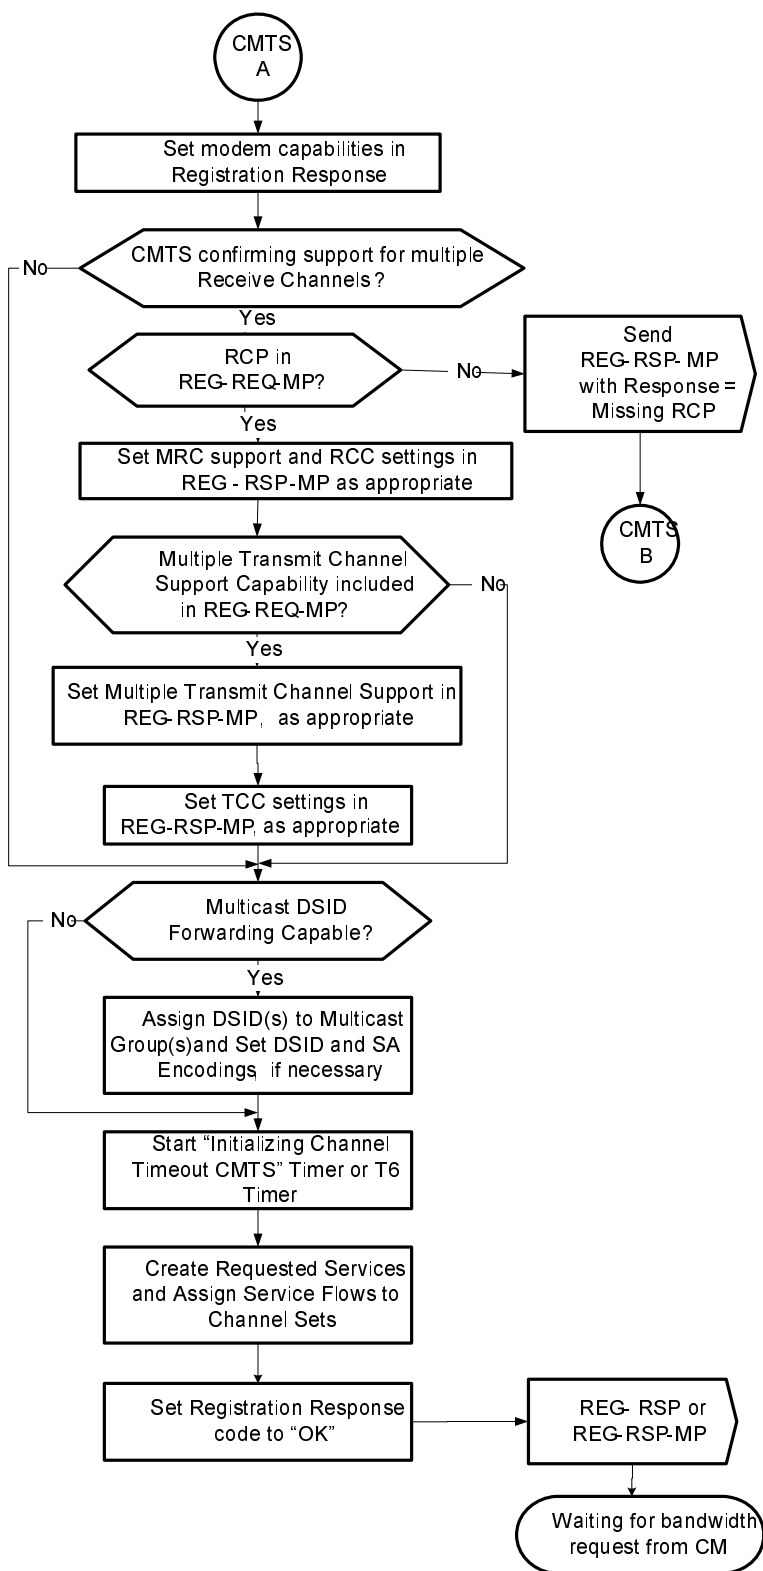


Figure 10-32: CMTS Registration - Continued

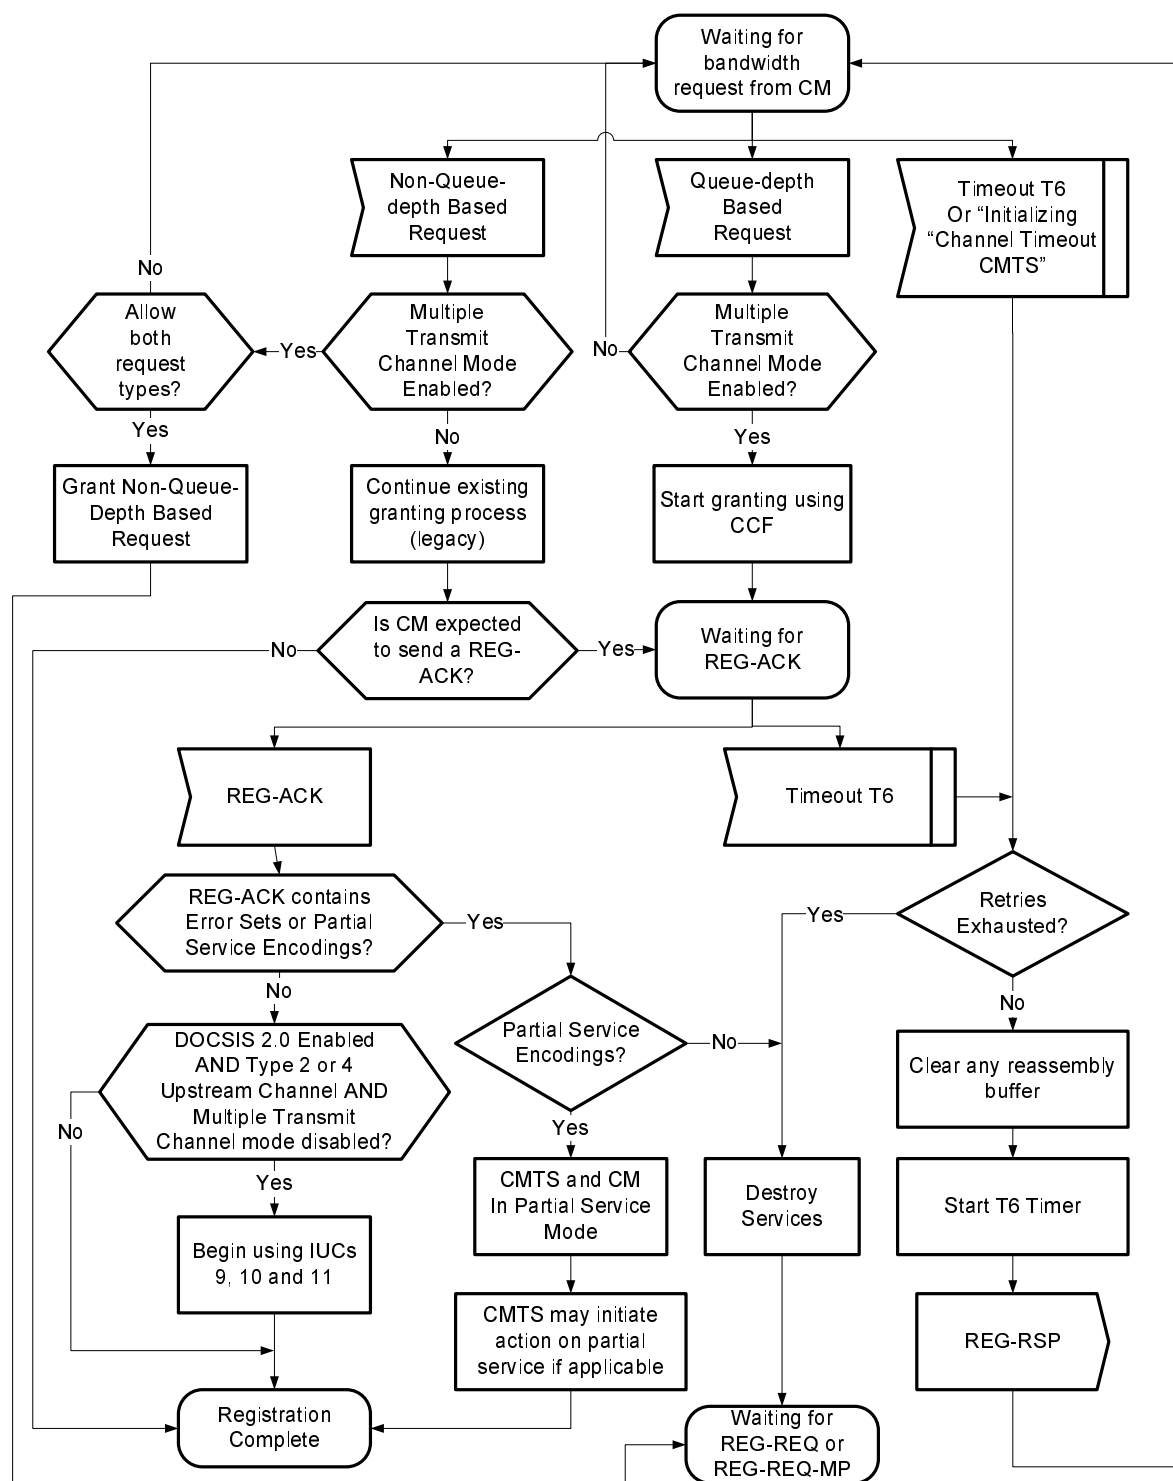


Figure 10-33: CMTS Registration - End

10.2.6.2.1 Channel Assignment During Registration

The Registration Request message can have a number of TLVs which will influence the selection of downstream and upstream channels that the CMTS assigns to CMs operating in Multiple Receive Channel (MRC) mode in the Registration Response. Additionally, some of these TLVs will cause the CMTS to re-direct CMs not operating in MRC mode to an alternate channel pair after Registration completes (via DCC or CM reboot).

To avoid potential conflicts between these TLVs there is a defined precedence order for handling of them by the CMTS. The CMTS MUST follow this precedence order for CMs operating in MRC mode:

- 1) TLV 56 Channel Assignment (clause C.1.1.25) and TLV 48 RCPs (clause C.1.5.3).

As described in clause 8.2.4, the CMTS is required to select an RCP from those advertised by the CM and generate a corresponding RCC that configures the CM's receivers. When the Channel Assignment TLV is present, the CMTS is additionally required (see clause C.1.1.25) to select an RCS and TCS that match the channel(s) indicated in the Channel Assignment TLV. If either or both of these cannot be achieved, the CMTS is required to reject the registration of the CM.

- 2) TLV 43.11 Service Type Identifier (clause C.1.1.18.1.10).

When the Service Type Identifier is present in the configuration file, the CMTS is required to select an RCS and TCS from a Restricted Load Balancing Group or MAC Domain that is available to the CM and that offers the signaled Service Type, if such RCS and TCS exist. If an RCS and/or TCS do not exist that provide the signaled Service Type the CMTS can assign an RCS and/or TCS that do not offer the signaled Service Type.

- 3) TLV 43.3 Load Balancing Group ID (clause C.1.1.18.1.3).

If the Service Type Identifier is not present in the Registration Request, the CMTS examines the Load Balancing Group ID, if present. If the available choices for RCS and TCS include channels associated with the signaled Load Balancing Group, the CMTS is required to assign the CM to the signaled Load Balancing Group. If these conditions cannot be met, the CMTS can disregard the Load Balancing Group ID. If the Load Balancing Group ID and Service Type Identifier both appear in the same configuration file, the CMTS is free to ignore the Load Balancing Group ID.

- 4) TLVs 24/25.31-33 Service Flow Attribute Masks (clauses C.2.2.3.6 to C.2.2.3.8).

If there are multiple RCSs and/or TCSs available that meet the requirements of the Service Type Identifier (if present) and Load Balancing Group ID (if present), the CMTS is required to select an RCS and/or TCS that meet the Required and Forbidden Attribute Masks of the Service Flows requested in the configuration file. If an RCS and/or TCS are not available that meet these criteria, the CMTS is free to choose an alternative RCS and/or TCS from among those previously identified.

- 5) TLV 43.9 CM Attribute Masks (clause C.1.1.18.1.8).

If there are multiple RCSs and/or TCSs available that meet the requirements of the Service Type Identifier (if present), Load Balancing Group ID (if present) and Service Flow Attribute Masks (if present), the CMTS can select an RCS and/or TCS that meet the CM Required and Forbidden Attribute Masks requested in the configuration file.

The CMTS MUST follow this precedence order for CMs not operating in MRC mode:

- 1) TLV 43.11 Service Type Identifier (clause C.1.1.18.1.10).

When the Service Type Identifier is present in the configuration file, the CMTS is required to select an upstream and downstream channel from a Restricted Load Balancing Group or MAC Domain that is available to the CM and that offers the signaled Service Type, if such channels exist. If an upstream and downstream channel do not exist that provide the signaled Service Type the CMTS can assign an upstream and/or downstream channel that do not offer the signaled Service Type.

- 2) TLV 43.3 Load Balancing Group ID (clause C.1.1.18.1.3).

After meeting the requirements for Service Type Identifier, the CMTS examines the Load Balancing Group ID, if present. If the available choices for upstream and downstream channel include a channel pair associated with the signaled Load Balancing Group, the CMTS is required to assign the CM to the signaled Load Balancing Group. If these conditions cannot be met, the CMTS can disregard the Load Balancing Group ID.

3) TLV 43.9 CM Attribute Masks (clause C.1.1.19.5).

If there are multiple upstream and/or downstream channels available that meet the requirements of the Service Type Identifier (if present) and Load Balancing Group ID (if present), the CMTS is required to select an upstream and/or downstream channel that meet the CM Required and Forbidden Attribute Masks requested in the configuration file. If an upstream and/or downstream channel are not available that meet these criteria, the CMTS is free to choose an alternative upstream and/or downstream channel. If an upstream and/or downstream channel are not available that meet these criteria, the CMTS can disregard the CM Attribute Masks.

4) TLVs 24/25.31-33 Service Flow Attribute Masks (clauses C.2.2.3.6 to C.2.2.3.8).

If there are multiple upstream and/or downstream channels available that meet the requirements of the Service Type Identifier (if present), Load Balancing Group ID (if present) and CM Attribute Masks (if present), the CMTS is required to select an upstream and/or downstream channel that meet the Service Flow Required and Forbidden Attribute Masks requested in the configuration file. If an upstream and/or downstream channel are not available that meet these criteria, the CMTS is free to choose an alternative upstream and/or downstream channel from among those already identified.

Note that the operator may configure the CMTS (via network management mechanisms) to restrict a particular CM to a certain Service Type ID and/or Restricted Group ID. If such a configuration is made, both the Service Type ID and Restricted Group ID configuration file TLVs are ignored by the CMTS.

If the TLVs present in the Registration Request message require the CMTS to move the CM to a different MAC Domain, the CMTS will need to force the CM to re-initialize in the new MAC Domain. While the exact mechanism is left to the vendor, the CMTS SHOULD minimize the time it takes for the CM to be redirected to the new MAC Domain. Examples of potential mechanisms are: the CMTS could allow the CM to complete registration (possibly with forwarding disabled) and then immediately send a DCC-REQ to the CM; the CMTS could send a Registration Response with a reject confirmation code or a RNG-RSP Abort, forcing the CM to re-initialize, then upon the subsequent ranging request, perform a downstream frequency override. In certain plant topologies, the CMTS may not be able to precisely determine a CM's initial location. This may occur when the CMTS has identified the MD-CM-SG of the CM in the original MAC Domain, but that MD-CM-SG identifies a set of fibre nodes rather than a single fibre node. In this situation, it may not be possible for the CMTS to identify the downstream frequency which will reach the CM in the desired new MAC Domain. The CMTS may need to make more than one attempt to direct the CM to the appropriate MAC Domain.

10.2.7 Baseline Privacy Initialization

Following registration, if the CM is provisioned to run Baseline Privacy and EAE was not enabled, the CM MUST initialize Baseline Privacy operations, as described in [15].

10.2.8 Service IDs During CM Initialization

After completion of the Registration process (clause 10.2.6), the CM will have been assigned Service IDs (SIDs) to match its provisioning. However, the CM needs to 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 MUST inform the CM of this assignment in the Ranging Response. The CMTS MAY monitor use of this SID and restrict traffic to that needed for initialization.

The CMTS MUST assign a temporary SID from the unicast SID space (see clause 7.2.1.2, Information Elements), for any CM that did not begin the initial ranging process with a B-INIT-RNG-REQ message. Any CM that began the initial ranging process with a B-INIT-RNG-REQ message is known at the time of initial ranging to support the expanded SID space and the CMTS MAY assign the CM a temporary SID from the expanded SID space. CMs MUST support the capability to transmit unicast traffic on the expanded SID space (see clause C.1.3.1.15). If a CM supports the above capability the CMTS MAY assign SID numbers from the expanded unicast SID space in the Registration Response.

Upon 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.

Upon 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 obtain a new temporary SID via the Upstream Channel Adjustment TLV or via initial ranging.

It is possible that the Ranging Response may be lost after transmission by the CMTS. The CM recovers 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-send the temporary SID that had previously been assigned to this CM. If the CMTS instead assigns a different temporary SID to this CM, the CMTS MUST make some provision for aging out the original temporary SID that went unused.

When assigning SIDs to provisioned Service Flows during registration, 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, 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.

10.3 Periodic Maintenance

Remote RF signal level adjustment at the CM is performed through a periodic ranging function using the RNG-REQ and RNG-RSP MAC messages. This is similar to initial ranging and is shown in figures 10-34 and 10-35. Note that each figure represents the operation for a single upstream channel on a CM.

The CMTS MUST provide each upstream channel in the CM's TCS 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 so that a MAP could be missed without the CM timing out. The size of this "subinterval" is CMTS dependent.

A CM which is not in Multiple Transmit Channel Mode MUST reinitialize its MAC with a CM Initialization Reason of T4_EXPIRED after T4 seconds have elapsed without receiving a Periodic Ranging opportunity.

When a CM is in Multiple Transmit Channel Mode and an upstream channel in the TCS incurs a T4 timeout or range request retries are exceeded, then that upstream channel is considered unusable for request or data transmissions and the CM enters a partial service mode in the upstream (see clause 8.4).

If all the upstream channels associated with the Primary Upstream Service Flow are unusable the CM MUST reinitialize its MAC with a CM Initialization Reason of NO_PRIM_SF_USCHAN. This is true even when there is a viable upstream remaining in the TCS and that upstream is not associated with the Primary Upstream Service Flow.

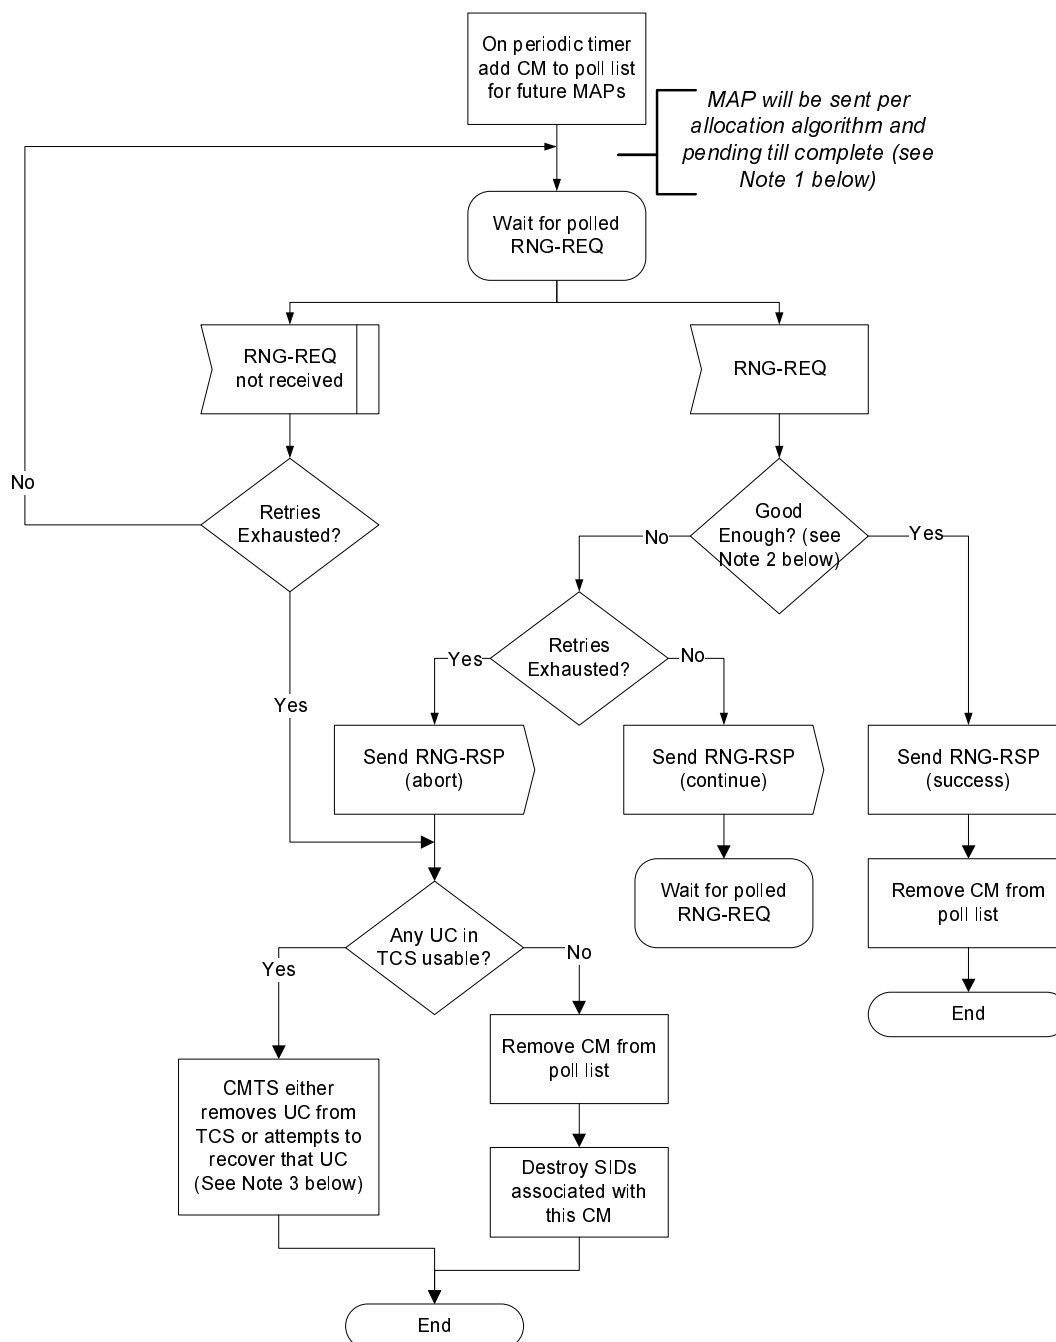
If there is at least one usable upstream channel associated with the Primary Upstream Service Flow, the CM MUST NOT reinitialize its MAC due to the loss of the other channels.

Upon receiving a RNG-RSP, the CM MUST use the first usable Reconfiguration Time or Global Reconfiguration Time (in the case of a dynamic range window adjustment), to adjust its transmitter parameters in accordance with the RNG-RSP.

To ensure interoperability with DOCSIS 1.0 and 1.1 CMTSs, both Multiple Transmit Channel Mode is disabled and not operating in DOCSIS 2.0 mode, the CM continues normal operation and MUST NOT suspend on-going data transmission if it receives a RNG-RSP with a ranging status of CONTINUE. When either Multiple Transmit Channel Mode is enabled or operating in DOCSIS 2.0 mode, the CM MUST NOT transmit anything other than RNG-REQs on a particular upstream channel, if that upstream channel has been suspended by receiving a RNG-RSP with a ranging status of CONTINUE, until such time as it receives a RNG-RSP with a ranging status of SUCCESS for the upstream channel in question.

The CMTS SHOULD NOT send a ranging status of CONTINUE in a RNG-RSP unless the ranging parameters measured on the corresponding RNG-REQ are insufficient for the CMTS to guarantee proper reception of all burst types available to that CM. Additionally, upon sending a RNG-RSP with ranging status of CONTINUE, the CMTS SHOULD schedule another Periodic Ranging opportunity for the CM on that upstream channel quickly so that the CM can return to a ranging status of SUCCESS as quickly as possible.

As described in clause 10.4.1, during normal operation in the S-CDMA mode, if a CM temporarily loses synchronization to the downstream signal, it is required to perform a ranging process before returning to normal operation. To facilitate this recovery, if the CMTS does not receive a RNG-REQ message from a CM during a Station Maintenance interval, the CMTS MAY schedule unicast Initial Maintenance opportunities or temporarily reduce the time between unicast spreader-off Station Maintenance opportunities.



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

NOTE 2: "Good Enough" means Ranging Request is within the tolerance limits of the CMTS for power, timing, frequency and transmit equalization (if supported).

NOTE 3: If the CMTS determines that there are still usable upstream channels in the CM's TCS associated with the Primary Service Flow, then it can attempt to recover using any method at its disposal. For example, the CMTS could re-range the unusable upstream channel by providing unicast maintenance opportunities; instruct the CM (via DBC-REQ) to alter its TCS to remove the unusable upstream channel; instruct the CM (via CM-CTRL-REQ message) to reinitialize its MAC.

Figure 10-34: Periodic Ranging - CMTS View

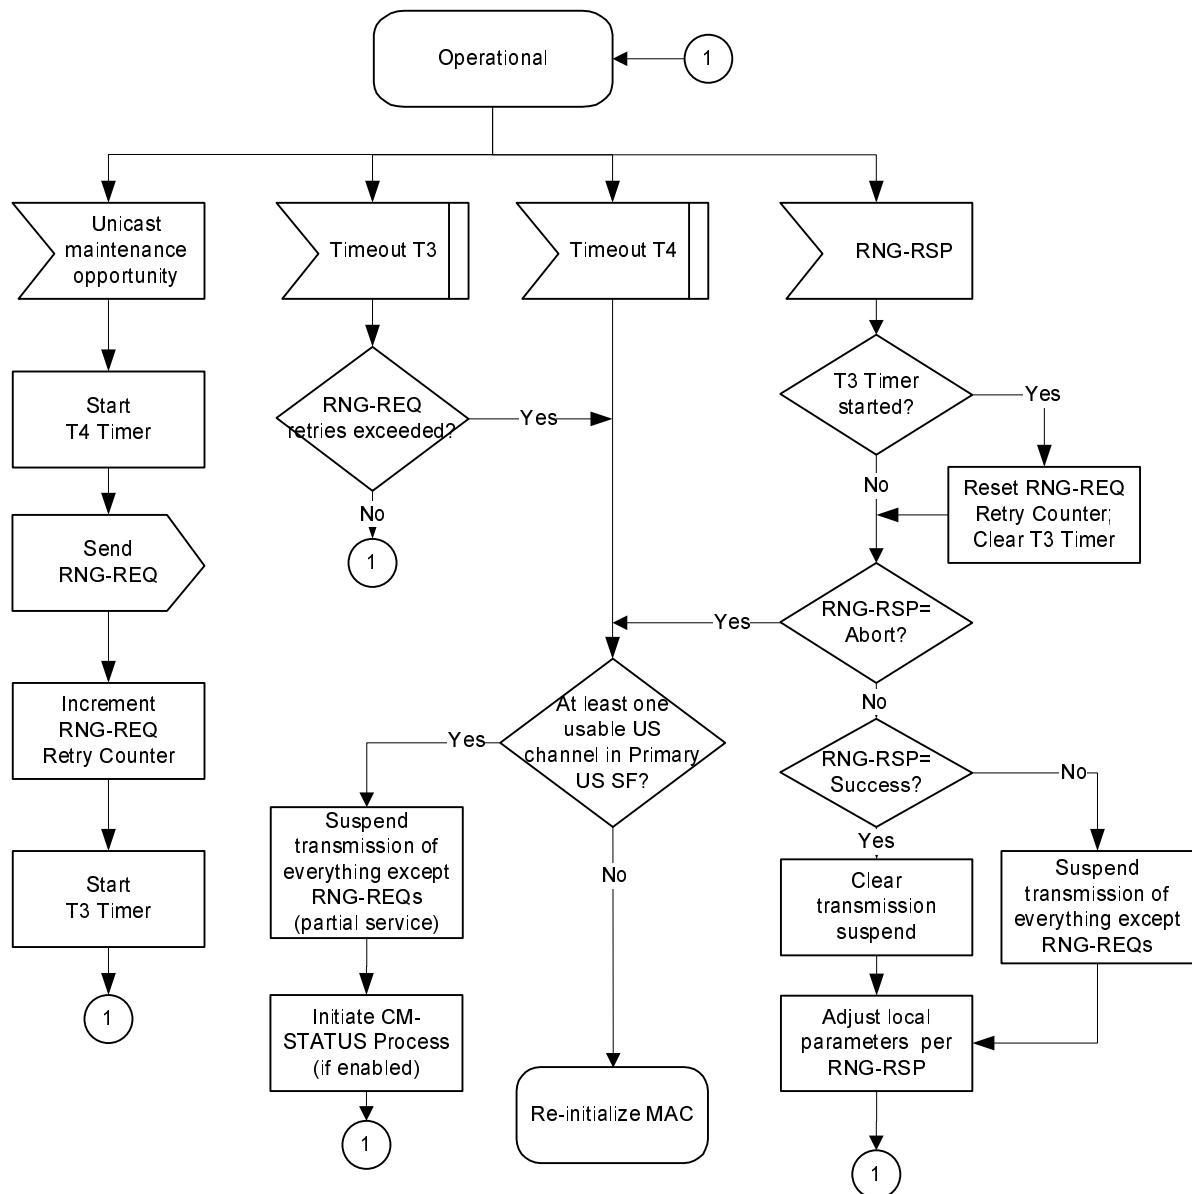


Figure 10-35: Periodic Ranging - CM View

10.4 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 [12] for details.
- At the Transmission Convergence layer, the CM can use the continuity counter and Payload Unit Start Indicator (PUSI) to detect and recover from lost MPEG packets [4].
- 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 clause 10.4.2 for details.
- All MAC management messages are protected with a CRC covering the entire message, as defined in clause 6. The CMTS MUST discard any message with a bad CRC. The CM MUST discard any message with a bad CRC. Table 10-2 shows the recovery process taken following the loss of a specific type of MAC message.
- At the network layer and above, the MAC Sublayer considers messages to be data packets protected by the CRC field of the data packet; any packets with bad CRCs are discarded. Recovery from these lost packets is in accordance with the upper layer protocol.

[10] contains a list of error codes with more useful information as to the failure of the PHY and MAC layers. Refer to clause 10.4.2 for additional information.

Table 10-2: Recovery Process on Loss of Specific MAC Messages

Message Name	Action Following Message Loss
SYNC	The CM can lose SYNC messages on the primary downstream for a period of the Lost SYNC interval (see annex B) before it has lost synchronization with the network. If the Lost SYNC Interval has elapsed without a valid SYNC message, the CM MUST suspend use of all upstream channels and try to re-establish synchronization again as described in clause 10.4.1.
MDD	Prior to registration, the CM uses the presence or absence of the MDD message to determine the appropriate initialization sequence as described in clause 10.4.2.1. After registration, the absence of an MDD message on a non-primary channel will be reported by the CM in a CM-STATUS message as specified in clause 6.4.34.
UCD	During CM initialization the CM has to receive a usable UCD before transmitting on the upstream channel. When in the "Collect UCDs" or "Obtain Upstream Parameters" state of the CM initialization process, if the CM does not receive a usable UCD within the T1 timeout period, the CM will continue scanning for a usable downstream channel. After having received a usable UCD for an upstream channel, whenever the CM receives a MAP with a UCD Count for that upstream channel that does not match the Configuration Change Count of the last UCD received, the CM suspends use of the corresponding upstream and begins looking for all UCD types for this upstream.
MAP	A CM is not allowed to transmit on an upstream channel without a valid upstream bandwidth allocation. If a MAP is missed due to error, the CM is not allowed to transmit on the corresponding channel for the period covered by the MAP.
RNG-RSP	If a CM fails to receive a valid ranging response within a defined time out period (T3) after transmitting a request, the CM retries the request a number of times defined in annex B as specified in clause 10.2.3.7. Failure to receive a valid ranging response after the requisite number of attempts causes the modem to declare the channel unusable as specified in clause 10.2.3.7.
REG-RSP	If a CM fails to receive a valid registration response within a defined time out period (T6) after transmitting a request, the CM retries the request a number of times defined in annex B as specified in clause 10.2.6. Failure to receive a valid registration response after the requisite number of attempts causes the modem to reinitialize MAC with a CM Initialization Reason of T6_EXPIRED as specified in clause 10.2.6.

10.4.1 CM Downstream Channel Interruptions

A Primary Downstream Channel 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.

In order to support redundant CMTS architectures, when a CM in the Operational state detects that the Primary Downstream Channel signal is invalid (i.e. does not meet the four criteria above), the CM MUST NOT immediately perform a Reinitialize MAC operation. The CM MUST instead attempt to re-establish synchronization on the current Primary Downstream Channel (see clause 7.3.1). The CM MUST attempt to re-establish synchronization until the operation of Periodic Ranging, as specified in figure 10-35, calls for a "Reinitialize MAC" operation after the expiration of Timeout T4 (with a CM Initialization Reason of T4_EXPIRED) or 16 expirations of Timeout T3 on all upstream channels in the CM's TCS (with a CM Initialization Reason of ALL_US_UNUSABLE).

An interruption of Primary Downstream Channel signal occurs when the following conditions are met:

- The interruption occurs on a downstream that is valid (per [12] and [14]) before and after the loss.
- The interruption is defined as an instantaneous loss of signal and after a predetermined delay, an instantaneous return to the original signal fidelity.
- The restored downstream signal is the original signal transmitted from the original source.

- The carrier frequency, physical plant and path delays remain the same before and after the interruption.
- There are no changes in any downstream signalling parameter, including the modulation and the M/N ratio, from before to after the interruption.

When a CM in the Operational state receives an interruption of Primary Downstream Channel signal less than or equal to 5 ms, the CM MUST recover from the outage such that its fixed timing error on S-CDMA channels is not greater than 2 % of the nominal modulation interval (in addition to the allowed jitter defined in [12]). When a CM in the Operational state receives an interruption of Primary Downstream Channel signal less than or equal to 5 ms, the CM MUST recover from the outage such that the first upstream transmission on TDMA channels after the CM resumes normal operation is performed within an accuracy of 250 nanoseconds plus 0,5 symbols (refer to [12]). On all upstream channels, the CM MUST continue with normal operation within 2 s from the end of the interruption. The CM is not required to continue normal operation if it receives a second interruption of downstream signal prior to the first receipt of a RNG-RSP with status "success".

When a CM in the Operational state receives an interruption of Primary Downstream Channel signal greater than 5 msec but less than the Lost Sync Interval (see annex B), the CM MAY continue with normal operation as long as it recovers within 2 s:

- With a fixed timing error not greater than 2 % of the nominal modulation interval (in addition to the allowed jitter defined in [12]) on S-CDMA channels.
- Within the accuracy specified in [12] on TDMA channels.

If the CM cannot recover according to the preceding recovery time, timing and jitter specifications, the CM MUST re-acquire upstream timing to an accuracy of at least 1 usec, be ready to respond to a ranging opportunity within 2 s and receive a RNG-RSP message with status "success" for a particular channel before resuming its upstream transmission on that channel. For the ranging process, the CM MUST use Broadcast or Unicast Initial Maintenance intervals or Station Maintenance intervals. However, a CM MUST NOT use spreader-on Station Maintenance on S-CDMA channels. For the ranging process, the CM MUST use the appropriate Ranging SID in the INIT-RNG-REQ or RNG-REQ message and use its known timing offset. A CMTS MUST process INIT-RNG-REQ messages with a Ranging SID from any CM that is in normal operation. If the Ranging SID used by the CM in INIT-RNG-REQ is no longer valid, the CMTS SHOULD send a RNG-RSP message to the CM with the ranging status set to "abort".

In all cases, after the first successful ranging opportunity subsequent to the interruption, the CM MUST meet the timing requirements specified in [12].

If CMTS becomes aware of an interruption of a CM's Primary Downstream Channel (via a CM-STATUS message from the affected CM or from another CM), it MAY send a DBC-REQ to the CM to select a new Primary Downstream Channel. If the CM receives the DBC-REQ prior to T4 timeout on all upstream channels, it will acquire SYNC on the new primary downstream channel, re-range on all upstreams and resume normal operation. This can prevent a significant interruption in service in the case of a primary downstream channel failure.

10.4.2 MAC Layer Error-Handling

This clause describes the procedures that are required when an error 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 HFC 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 type of error, 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.

The CM MUST verify the HCS of every received MAC Frame. When a bad HCS is detected, the CM MUST discard the MAC Header and any payload. The CMTS MUST verify the HCS of every received MAC Frame. When a bad HCS is detected, the CMTS MUST discard the MAC Header and any payload.

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. The CMTS MUST verify the CRC of every received Packet PDU or Isolation Packet PDU MAC Frame. When a bad CRC is detected, the CMTS MUST discard the PDU portion of the Packet PDU or Isolation Packet PDU MAC Frame. The CM MUST verify the CRC of every received Packet PDU or Isolation Packet PDU MAC Frame. When a bad CRC is detected, the CM MUST discard the PDU portion of the Packet PDU or Isolation Packet PDU MAC Frame.

Requirements for reporting of Error Codes and Messages by the CM and CMTS are described in [10].

10.4.2.1 Error Recovery During Pre-3.0 DOCSIS Fragmentation

There are some special error handling considerations for fragmentation. Each fragment has its own fragmentation header complete with a Fragment Header Checksum (FHCS) and its own FCRC. There may be other MAC headers and CRCs within the fragmented payload. However, only the FHCS and the FCRC are used for error detection during fragment reassembly.

If the FHCS fails the CMTS MUST discard that fragment. If the FHCS passes but the FCRC fails, the CMTS MUST discard that fragment. The CMTS MAY process any requests in the fragment header of a fragment that was discarded for an FCRC failure. The CMTS SHOULD process such a request if it is performing fragmentation in Piggyback Mode (refer to clause 7.2.5.2.2.2). This allows the remainder of the frame to be transmitted by the CM as quickly as possible.

If a CMTS is performing fragmentation in Multiple Grant Mode (refer to clause 7.2.5.2.2.1), it SHOULD complete all the grants necessary to fulfill the CM's original request even if a fragment is lost or discarded. This allows the remainder of the frame to be transmitted by the CM 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 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).

10.4.2.2 Error Recovery During Segmentation with Segment Headers On

There are some special error handling considerations for segmentation with Segment Headers On. Each segment has its own segment header complete with an HCS. If the HCS for a segment fails, the CMTS MUST discard that segment. If the HCS passes for a segment, the CMTS may process any bandwidth request in the segment header prior to reordering the segments and reassembling the received packet stream.

The CMTS uses the sequence number in the segment header to know the order of the segment relative to other segments for that service flow. Once the CMTS receives a higher sequence number on each of the active upstream channels associated with a service flow, the CMTS knows that any missing lower sequence numbers have been lost. Once the CMTS has placed the received segments in the proper order, it uses the pointer field in the segment headers to find the first MAC frame header (if present) in the segment. The CMTS uses the length fields in the DOCSIS headers along with the HCS to determine if the DOCSIS Header or packet payload is spanning the segment boundary. Once the packet payload is identified, the CRC is verified.

Should the HCS in a packet header within a segment fail, the CMTS MAY discard the remainder of that segment and begin processing with the next DOCSIS header in a subsequent segment. The CMTS MUST discard any partial packets during this process if the remaining pieces cannot be determined. The CMTS MUST forward any complete packets in the correct order according to the sequence number in the segment headers.

In addition, the CMTS MAY restart the segment reassembly process based on implementation dependent criteria such as a reassembly timer.

10.4.3 CM Status Report

CM-STATUS messages are needed in cases where the CM detects a failure that the CMTS can not detect directly (for example, a failure in the CIN where an M-CMTS is used) or where the CM can send valuable information to the CMTS when an error or a recovery event occurs (for example, the CM can report a T3 timeout to the CMTS). Upon receiving an error indication the CMTS is expected to take action in order to correct the error.

A CM MUST transmit a CM-STATUS message on any available channel when it detects an event condition listed in table 10-3 for any object and reporting of the event type for that object is enabled on the CM. Table 10-3 describes the trigger conditions that set each event "on" and the reset conditions at which the event is considered to change to "off". An event is said to "occur" when it transitions from "off" to "on".

Some event types are for a particular downstream channel, a particular upstream channel or a DSID. For each such event, the CM maintains a separate state variable as to whether the event condition is considered "on" or "off" for each channel or DSID.

The CM MUST NOT send a CM-STATUS message if the CM-STATUS Event Control TLV (see clause 6.4.28.1.11) in the MDD message is not specified for a particular event type. An event type cannot be enabled until the CM-STATUS Event Control TLV for the event is specified in a subsequent MDD message.

If the CM-STATUS Event Control TLV in the MDD message is specified for a particular event type, the CM MUST enable/disable event reporting for channel specific events according to the following:

- 1) If an Override for Status Event Enable Bitmask for a channel is specified via a unicast CM--CTRL-REQ message, then the CM enables/disables event reporting for the event type on the channel according to the bitmask specified in the CM-CTRL-REQ message.
- 2) If an Override for Status Event Enable Bitmask is not specified via a unicast CM-CTRL-REQ message, the CM discards any previously received Override for that channel and reverts back to the CM-STATUS Event Enable Bitmask provided in the MDD message.

If the CM-STATUS Event Control TLV in the MDD message is specified for a particular event type, the CM MUST enable/disable event reporting for non-channel specific events according to the following:

- 1) If an Override for the CM-STATUS Event Enable Bitmask for Non-Channel-Specific Events is specified via a unicast CM-CTRL-REQ message, then the CM enables/disables the event reporting for the event type according to the bitmask specified in the CM-CTRL-REQ message.
- 2) If an Override for the CM-STATUS Event Enable Bitmask for Non-Channel-Specific Events is not specified via a unicast CM-CTRL-REQ message, the CM discards any previously received Override for the CM-STATUS Event Enable Bitmask for Non-Channel-Specific Events and reverts back to the CM-STATUS Event Enable Bitmask for Non-Channel-Specific Events provided in the MDD message.

The CM MUST NOT send a CM-STATUS message for an event type for which the reporting has been disabled. The CM MUST NOT send a CM-STATUS message prior to becoming operational.

A Primary Channel MDD Timeout event is said to occur if the Lost MDD Timeout has passed without receipt of a valid MDD message on the CM's Primary Downstream Channel. During a Primary Channel MDD Timeout event, the CMTS is unable to control CM-STATUS reporting of the affected CM. Therefore, during a Primary Channel MDD Timeout event, the CM MUST NOT send CM-STATUS messages. The CM MUST disable any event reporting and reset the state machines to IDLE. Upon receipt of a valid MDD message following a Primary Channel MDD Timeout event, the CM MUST re-process the Primary MDD message and re-enable event reporting according to the new primary MDD.

When one or more events of the same event type are "on" and enabled for reporting, the CM sends a CM-STATUS message that reports the event condition for all such events.

For each event type, the CM maintains the following state information:

- A Transaction Identifier that identifies each uniquely reported transition of one or more events of the event type from off to on.
- A Maximum Holdoff Timer value that controls how often repeated CM-STATUS messages for the same Transaction Identifier are sent.
- A Maximum Reports Count that controls how many CM-STATUS messages for the same Transaction Identifier are transmitted by the CM. A Maximum Reports Count of zero signals that the CM continues sending CM-STATUS messages as long as the event condition is "on" and is enabled for reporting.
- A "ReportsLeft" counter of the number of reports of an event type's Transaction Identifier left to be reported to the CMTS.

The CM updates its Maximum Holdoff Timer and Maximum Reports Count for an event type when the CM-STATUS Event Control Encoding in the CM's primary channel MDD changes these values.

For each event type, the CM MUST maintain a CM-STATUS Process State Machine described by the SDL description below that controls the timing and number of CM-STATUS report messages sent by the CM for the event type. Each CM-STATUS message reports a single event type condition for all relevant, downstream channels, upstream channels or DSIDs. For the "Sequence Out of Range" event type for DSIDs, the Maximum Holdoff Timer can be overridden for an individual DSID by the CMTS (clause C.1.5.4.3.5). In this case, the CM implements a separate CM-STATUS Process State Machine for each event corresponding to the Sequence Out of Range event type for a DSID with an overridden Maximum Holdoff Timer.

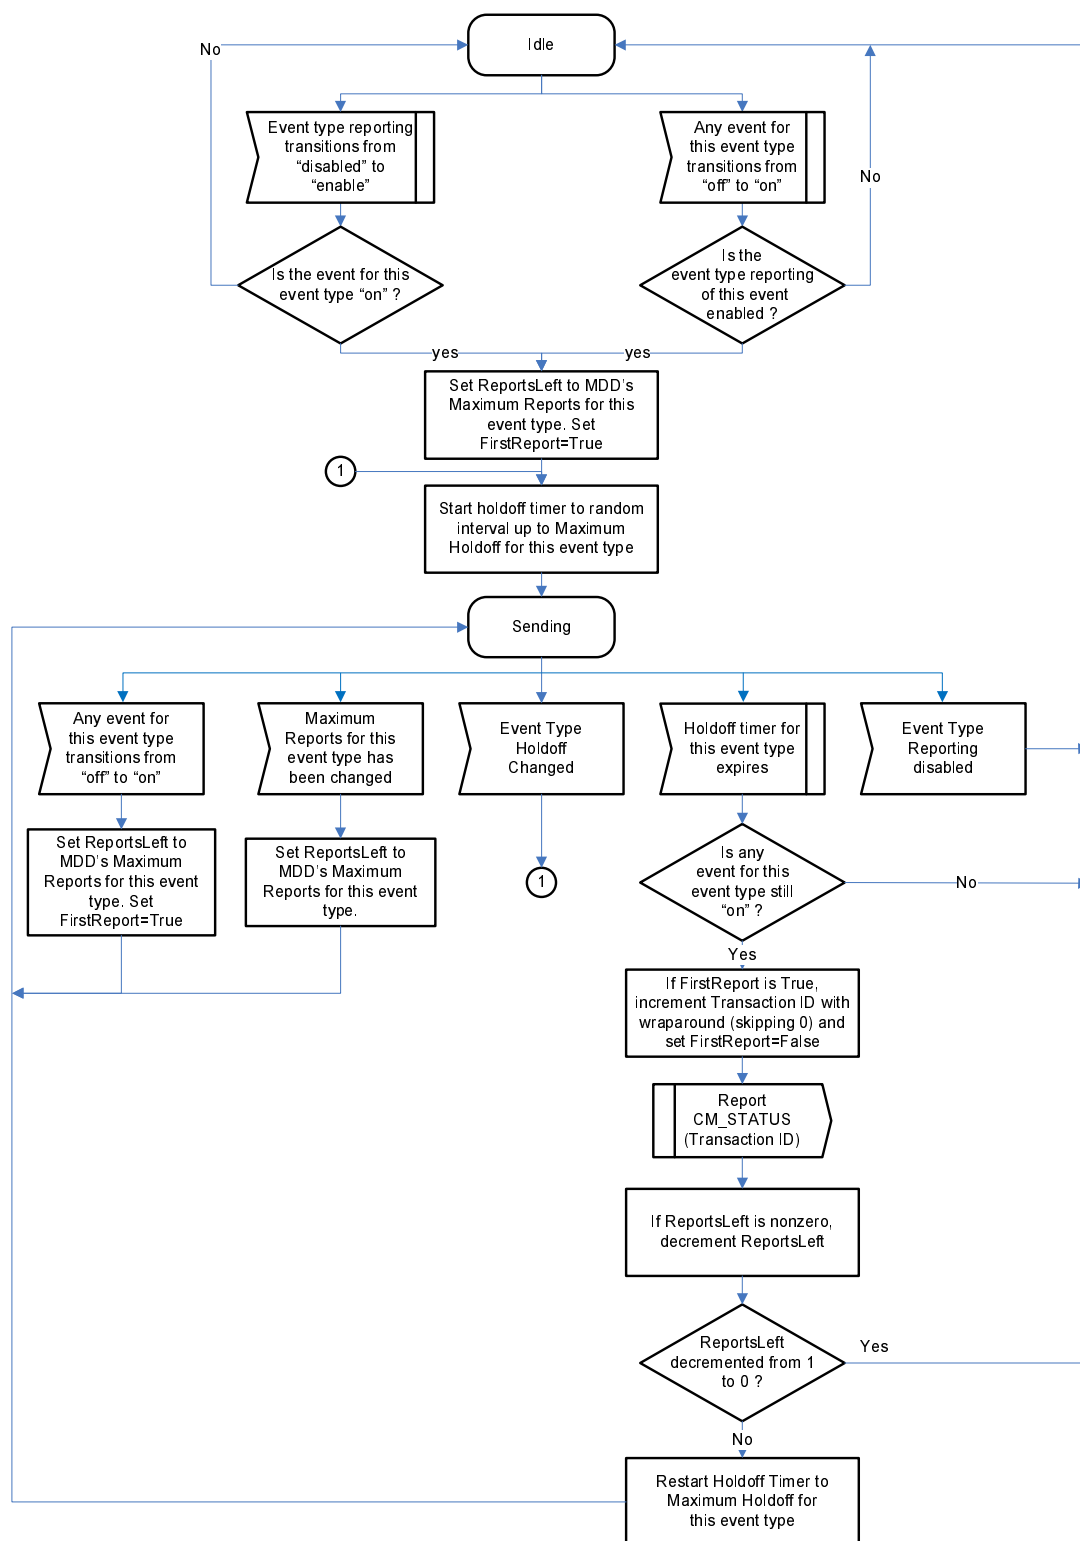


Figure 10-36: CM-STATUS Event Type State Machine

Operation of the CM-STATUS Event Type State Machine is described below.

The CM is considered to be in two states for each CM-STATUS event type: IDLE and SENDING. The state machine starts in the IDLE state with the Transaction Identifier variable set to 0.

When an event occurs (i.e. transitions to "on") in the IDLE state and the event type reporting for that event is enabled or when the event type reporting transitions from "disabled" to "enable" in the IDLE state and an event for this event type has been "on", the CM sets the ReportsLeft variable to the Maximum Reports setting for the event type and sets the "FirstReport" control variable to True. The CM then selects an initial report holdoff timer value randomly between 0 and the value specified by the Maximum Holdoff for the event type. The granularity of the holdoff timer should be as fine as possible, but no less than 20 milliseconds. The CM then enters the SENDING state for the event type and remains in the SENDING state whenever the holdoff timer is running. The initial choice of a random holdoff timer interval is intended to prevent the flooding of CM-STATUS messages in cases where a failure has affected a large number of CMs.

When the holdoff timer expires in the SENDING state, the CM first verifies that at least one event for the event type remains on. If not, the CM returns the event type to the IDLE state without sending a CM-STATUS message and possibly without having incremented the Transaction Identifier for the event type. If any event remains "on", the CM checks whether the CM-STATUS message it is about to send is the first report of a new transaction. If so, the CM clears its "FirstReport" control flag and increments the Transaction Identifier for the first CM-STATUS report of a new Transaction. The CM wraps the 16-bit Transaction Identifier from 65535 back to 1, skipping 0 when it wraps around. If it is not the first report of a new report transaction, the CM leaves the Transaction Identifier variable for the event type unchanged. The CM then sends the CM-STATUS message, including separate Event Encoding TLVs for each enabled event of the event type.

After the CM transmits the CM-STATUS message, it checks whether the ReportsLeft variable is already zero, indicating that Maximum Reports for the event type was also zero, which means that reports are sent until disabled. If Reports left was not already zero, the CM decrements the ReportsLeft variable for the event type. If it decrements ReportsLeft from one to zero in this case, all CM-STATUS messages for a transaction have been sent and the CM returns to the IDLE state for the event type. Otherwise, i.e. when an additional CM-STATUS report for the transaction is required, the CM re-starts the holdoff timer to the Maximum Holdoff Timer value for the event type and returns to the SENDING state. Thus, for a single event type transaction reported from the IDLE state, the first CM-STATUS report is sent with a random holdoff timer and all subsequent reports from the SENDING state are sent with the fixed, maximum timer for the event type.

NOTE: Other events of the same event type may turn "on" while awaiting the sending of a CM-STATUS report for an original event that causes the IDLE to SENDING transition. Furthermore, the original event may turn "off" and then back "on" while awaiting the sending of the first CM-STATUS message. When any event of the event type transitions from "off" to "on" while in the sending state, the CM sets the ReportsLeft counter for the event type back to its Maximum Reports value and sets the FirstReport flag to True. When the current holdoff timer for the SENDING state expires, this will cause the CM to increment the Transaction Identifier.

While the CM is in the SENDING state for a particular event type, if the CMTS disables CM-STATUS reporting for the event type, the CM transitions to the IDLE state.

If the CM detects in its primary MDD that the Maximum Holdoff for an event type has changed while it is in the SENDING state for that event, it recalculates its current holdoff timer to a random interval up to the new maximum holdoff value and resumes waiting for the new holdoff timer in the SENDING state. The ReportsLeft variable is not changed in this case.

Each CM-STATUS message contains event reports of a single event type code.

10.4.3.1 Event Codes

As described above, reporting for each of these events is controlled by CM-STATUS Event Enable Bitmask and CM-STATUS Event Control TLV in the MDD message and CM-CTRL-REQ message.

The CM power events (Codes 9 and 10) are only applicable to CMs with battery backup capability. These events are used to signal the CMTS when the CM is operating on battery power. If the CMTS receives a CM-STATUS message with "CM operating on battery backup" indicated, the CMTS SHOULD send DBC messaging to the CM to reduce the CM's operation to a single upstream and downstream channel. This is because the CM's battery life will be shortened while transmitting or receiving on multiple channels.

For more details see annex P. Table 10-3 lists the CM-STATUS message codes.

Table 10-3: CM-STATUS Event Type Codes and Status Events

Event Type Code	Event Condition	Status Report Events		Parameters Reported		
		Trigger Event to "on"	Reset Event to "off"	Downstream Channel ID	Upstream Channel ID	DSID
0	Reserved					
1	Secondary Channel MDD timeout	Lost MDD Timer expiry of a secondary channel advertised as active in the primary channel MDD.	Receipt of MDD; OR removal of the channel from the active channel list in the primary channel MDD; OR removal of the channel from the CM's Receive Channel Set via DBC-REQ.	The CM MUST report the Channel ID upon which the trigger event occurred.	N/A	N/A
2	QAM/FEC lock failure	Loss of QAM or FEC lock on one of the downstream channels advertised as active in the primary channel MDD.	Re-establishment of QAM/FEC lock; OR removal of the channel from the active channel list in the primary channel MDD; OR removal of the channel from the CM's Receive Channel Set via DBC-REQ.	The CM MUST report the Channel ID upon which the trigger event occurred.	N/A	N/A
3	Sequence out-of-range	Receipt of a packet with an out-of-range sequence number for a particular DSID.	Receipt of a packet with an in-range sequence number; OR change in the Sequence Change Count.	N/A	N/A	The CM MUST report the DSID upon which the trigger event occurred.
4	Secondary Channel MDD Recovery	Receipt of an MDD on a Secondary channel advertised as active in the most recent primary channel MDD.	MDD timeout event on the channel; OR removal of the channel from the active channel list in the primary channel MDD; OR removal of the channel from the CM's Receive Channel Set via DBC-REQ.	The CM MUST report the Channel ID upon which the trigger event occurred.	N/A	N/A

Event Type Code	Event Condition	Status Report Events		Parameters Reported		
		Trigger Event to "on"	Reset Event to "off"	Downstream Channel ID	Upstream Channel ID	DSID
5	QAM/FEC Lock Recovery	Successful QAM/FEC lock on a channel advertised as active in the most recent primary channel MDD.	Loss of QAM/FEC lock; OR removal of the channel from the active channel list in the primary channel MDD; OR removal of the channel from the CM's Receive Channel Set via DBC-REQ.	The CM MUST report the Channel ID upon which the trigger event occurred.	N/A	N/A
6	T4 timeout (note 1)	Expiration of the T4 timeout on the CM.	Receipt of maintenance opportunity (initial maintenance or station maintenance); OR removal of the channel from the active channel list in the primary channel MDD; OR removal of the channel from the CM's Transmit Channel Set via DBC-REQ.	N/A	The CM MUST report the Channel ID upon which the trigger event occurred.	N/A
7	T3 re-tries exceeded	The number of T3 retries as specified in annex B is exceeded.	Receipt of RNG-RSP message ; OR removal of the channel from the active channel list in the primary channel MDD; OR removal of the channel from the CM's Transmit Channel Set via DBC-REQ.	N/A	The CM MUST report the Channel ID upon which the trigger event occurred.	N/A

Event Type Code	Event Condition	Status Report Events		Parameters Reported		
		Trigger Event to "on"	Reset Event to "off"	Downstream Channel ID	Upstream Channel ID	DSID
8	Successful ranging after T3 re-tries exceeded	Successful ranging on a channel for which T3 re-tries exceeded event had been reported.	The number of T3 retries as specified in annex B is exceeded; OR removal of the channel from the active channel list in the primary channel MDD; OR removal of the channel from the CM's Transmit Channel Set via DBC-REQ.	N/A	The CM MUST report the Channel ID upon which the trigger event occurred.	N/A
9	CM operating on battery backup	CM detects loss of A/C Power for more than 5 s and the CM is operating on battery backup.	CM detects the presence of A/C Power and has returned from backup battery to operating on A/C power.	N/A	N/A	N/A
10	CM returned to A/C power	CM detects the presence of A/C Power for more than 5 s and has returned from backup battery to operating on A/C power.	CM detects loss of A/C Power and the CM is operating on battery backup.	N/A	N/A	N/A
11 to 255	Reserved for future use					

10.5 DOCSIS Path Verification

10.5.1 DPV Overview

The DOCSIS Path Verify (DPV) protocol offers two modes of operation:

- 1) **Per Path:** An operational mode which will permit the measurement of latency between two particular DPV reference points. This mode uses a dedicated MAC Management Message to perform the measurement.
- 2) **Per Packet:** A diagnostics mode where the source (either the CM or CMTS) will attach a diagnostic extended header to each packet within a specified service flow. This header is intended to be intercepted by external test equipment and ignored by the rest of the system.

Messages which are inserted per path can be done so independent of the existence of data packets within that path.

10.5.2 DPV Reference Points

The reference points recognized by DPV are shown in figure 10-37. The expression "DS MAC" refers to the downstream MAC processing element and the term "US MAC" refers to the upstream MAC processing element.




Figure 10-37: DPV Reference Diagram

Each direction, downstream and upstream, has a separate set of reference points. Table 10-4, DPV Downstream Reference Point Descriptions and table 10-5, DPV Upstream Reference Point Descriptions, provide a more precise description of each DPV reference point.

Table 10-4: DPV Downstream Reference Point Descriptions

Reference Point	Code Assignment	Description
	0	A value of 0 is reserved to indicate that a reference point is not being specified.
D1	1	A reference point in the CMTS that generally represents the input to the DOCSIS MAC. Note that the time between D1 and D2 includes time spent on a DOCSIS service flow queue including maximum rate limiting and QoS scheduling delays. This point is individually determined by the CMTS manufacturer.
D1'	2	A reference point in the CMTS that generally represents the output of the DOCSIS MAC. For an Integrated CMTS, this is prior to the R-S encoder and QAM modulator. This point is typically where SYNC insertion takes place and is generally a fixed delay (depending upon interleaver depth) from the actual RF output. For a Modular CMTS, it is located at the M-CMTS Core DEPI output. Note that M-CMTS Cores which employ internal data paths after the DOCSIS MAC circuitry may have additional latency which may become included in any measurement that starts at D2. This point is individually determined by the CMTS manufacturer.
D2	11	CM RF Interface. This point is located after the tuner, QAM demodulator, de-interleaver and R-S decoder, but prior to input packet queuing. The measurement point is with respect to the end of the received packet.

Table 10-5: DPV Upstream Reference Point Descriptions

Reference Point	Code Assignment	Description
	0	A value of 0 is reserved to indicate that a reference point is not being specified.
U1	21	A CM reference point that generally represents the input to the DOCSIS MAC processing element, before maximum rate limiting, QoS scheduling delays and Request-Grant latencies.
U2	31	A CMTS upstream receiver reference point that is individually determined by the CMTS manufacturer.

For the DPV Per Path operation with the DPV MAC Management Message, the CMTS MAY support DPV reference points D_1 , D_1' and U_2 . For the DPV Per Path operation with the DPV MAC Management Message, the CM MUST support downstream DPV measurements to reference point D_2 . For the DPV Per Path operation with the DPV MAC Management Message, the CM MAY support upstream DPV measurements from reference point U_1 . The CMTS SHOULD perform a measurement between reference points D_1 and D_2 . There are no requirements on the internal latency of the CM between the reception of a DPV-REQ and the generation of a DPV-RSP. Measurements that include the internal latency of the CM may be highly variable.

For the DPV per Packet operation with the DPV Extended header, the CM MAY support reference point U_1 .

For measurement point D_2 the CM MUST use a timestamp value derived from the downstream SYNC message that has not been adjusted by the CM ranging process. If the CM supports measurement point U_1 the CM MUST use a timestamp value derived from the downstream SYNC message that has been adjusted by the CM ranging process (i.e. the current upstream mini-slot timestamp). If the CM does not support upstream reference point U_1 , it MUST insert a timestamp value of 0 in any DPV-RSP which includes U_1 . The CM MUST insert a timestamp value that is within 1 ms of its actual current timestamp value. The CM SHOULD insert a timestamp value that is within 100 usec of its actual current timestamp value.

10.5.3 DPV Math

The difference between the Timestamp End and the Timestamp Start in the DPV-RSP (clause 6.4.33) does not include downstream propagation delay in the HFC and thus should be considered a relative latency rather than an absolute latency. The reason for this has to do with how timestamps are used and distributed in a DOCSIS system. The CMTS distributes a timestamp through the SYNC message to the CM. If a measurement packet was to travel the same path with the same latency as the SYNC message, with a start point in a CMTS and an endpoint in the CM, the resulting formula:

- Relative Latency = Timestamp End - Timestamp Start.

would result in a relative latency of zero, even though there obviously is latency in the HFC path. The observation is that because downstream latency measurements are in the same direction as the SYNC message, the measurement does not include the latency seen by the SYNC message. Note that use of the CM ranging offset does not solve the accuracy problem as the CM ranging offset may vary between CM manufacturers depending upon individual internal circuit delays.

There is an additional latency error the CMTS may want to compensate for. The CMTS will insert a timestamp into the DPV-REQ packet prior to transmission. The CM will insert a timestamp into the packet after the reception of the message. Thus, the delta of the two timestamps includes the serialization time of the packet. The serialization time is the time it takes to transmit the DPV packet onto the QAM Channel. This error also exists in the upstream direction.

The difference between any two relative latency measurements can be considered as a valid skew measurement. As such, skew can be measured between two flows within or across QAM Channels. This is intended to be useful for detecting congestion latency in an M-CMTS EQAM and determining its impact upon downstream resequencing.

There is no bound on the CM internal processing time between reception of the DPV-REQ message and the transmission of the DPV-RSP. As such, any round-trip latency measurement includes this implementation-specific (and possibly variable) processing time and cannot be used to accurately compare round trip times between devices.

When the CM needs to calculate the average latency, it uses a running average. If N is held constant, the type of running average in the formula is known as an Exponential Moving Average (EMA). An EMA places a heavier weight on more recent samples as opposed to a Simple Moving Average (SMA) which places an equal weight on all samples. The CM MUST use the following formula for its running average latency calculations:

- Average Latency' = Average Latency + Alpha * (Last Measured Latency - Average Latency);

where:

- Alpha = $1 / N$.

Average Latency' represents the updated value of Average Latency. The value of N is supplied in the DPV-REQ message. N can be dynamically chosen by the CMTS such that Alpha is a number between 0 and 1 and represents a weighting for the current sample, relative to the weight given to the accumulated average.

10.5.4 DPV Per Path Operation

The DPV Per Path feature is appropriate for sampling the latency of a particular data path and for generating long term averages. DPV Per Path measurements can be made independent of the data packet flow.

Latency measurements may be useful in the downstream direction for several applications, including the determination of the skew of a bonding group by comparing latency between different QAM Channels within the bonding group.

The DPV Per Path operation is achieved through the use of two unique MAC management messages. The first message, DPV-REQ is sent from the CMTS to the CM. The second message, DPV-RSP is sent from the CM to the CMTS. All measurements are originated by the CMTS. There is an Echo bit within the DPV-REQ header which indicates to the CM that it should generate a DPV-RSP.

When the CMTS wants to make a latency measurement, it generates a DPV-REQ MAC management message. The latency measurement is done between two reference points known as the start reference point and the end reference point. The start reference point may be any supported reference point in the downstream or upstream direction. The end reference point may be any supported reference point, but **MUST** be a point that occurs after the indicated start reference point.

For measurements that start and end in the downstream direction, the CM **MUST** maintain two independent sets of statistics per Downstream QAM Channel each of which reflect:

- **Last Measured Latency:** This contains the most recent latency measurement.
- **Minimum Latency:** This contains the lowest latency value measured since the last clearing of the DPV statistics.
- **Maximum Latency:** This contains the highest latency value measured since the last clearing of the DPV statistics.
- **Average Latency:** This contains a running average of the latency value over the entire history of measurements since the last clearing of the DPV statistics (see clause 10.5.3).

The two sets of statistics permit different downstream flows to be compared. The CMTS indicates in which statistics set a particular measurement should be included. The CMTS can also reset the statistics with the DPV-REQ message. These values **MUST** be readable through the CM MIB.

This allows the CMTS to pursue two different measurement techniques. The CMTS could send a measurement packet with the echo bit set and perform analysis at the CMTS on each measurement. Alternatively, the CMTS could send a series of measurement packets with the echo bit not set and have the CM perform the measurement analysis. The results could then be retrieved by the CMTS from the CM as needed.

10.5.4.1 DPV Ping

A specific usage of DPV Per Path Operation is known as a "DPV Ping". A DPV Ping consists of a DPV MAC message exchange with the Echo bit asserted and with the remaining parameter values of DPV-REQ and DPV-RSP cleared except the transaction ID.

10.5.5 DPV Per Packet Operation

The DPV Per Packet operation is appropriate for determining the maximum and minimum latency seen by the packets of a particular service flow. DPV Per Packet operation can only be performed when data packets are present in the service flow.

The DPV Per Packet operation is performed by having the source device generate and append a DPV Extended Header to each packet it transmits on a given service flow or flows. The receiving device is presumed to be a network sniffer or other diagnostic device. The CM DPV Per Packet operation is enabled and disabled through the CM MIB. The CMTS and CM are not required to perform any action upon the reception of the DPV extended header.

It should be noted that if the DPV extended header is enabled on a UGS flow in the upstream, that the UGS scheduling at the CMTS will have to be modified to accommodate the increased packet size. How this is achieved is outside the scope of the present document.

11 Dynamic Operations

11.1 Upstream Channel Descriptor Changes

Whenever the CMTS is to change any of the upstream burst characteristics specified in the Upstream Channel Descriptor (UCD) message (see clause 6.4.3), it needs to 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 characteristics, it **MUST** announce the new values in an UCD message and increment the Configuration Change Count field in that UCD message to indicate that a value has changed. However, the CMTS **MUST NOT** start the UCD change process on an US channel if one or more CMs using this channel are still handling a previously initiated management transaction (like previous UCD change, DBC, DCC, etc.) that involves this US.

After transmitting one or more UCD messages with the new change count value for each UCD type to be used for this US, the CMTS transmits a MAP message with a UCD Change Count matching the new Configuration Change Count. The first interval in this MAP **MUST** be a data grant to the null Service ID of at least 1,5 ms for a TDMA channel or for the longer of 1,5 ms or the duration of 2 S-CDMA frames for an S-CDMA channel (to allow for the latency of the S-CDMA framing). When the change affects an S-CDMA channel the Start Time of the MAP with the new UCD Change Count **MUST** correspond to the beginning of an S-CDMA frame.

The CMTS **MUST** allow this time for cable modems to change their PMD sublayer parameters to match the new set. This time is independent of the lead time the CMTS needed to allow for in transmitting the MAP (see clause 7.2.1.6). The CMTS **MUST** transmit the new UCD message early enough that the CM receives the UCD message at least the UCD Processing Time (see annex B) prior to the time the first MAP using the new UCD parameters arrives at the CM.

With the exception of the following cases the CM **MUST** be able to transmit normally on the first grant following the grant to the NULL SID:

- 1) When the new UCD message has changed the S-CDMA Enable parameter.
- 2) When the new UCD message has changed the S-CDMA US Ratio Numerator or Denominator.
- 3) When UCD changes for multiple upstream channels within the TCS take effect within 1,5 ms of each other as described by the MAP messages.

In the first two cases the CM **MAY** redo initial ranging to establish timing synchronization for the new mode of operation before it resumes normal transmissions. If the CM was already registered with the CMTS and it redoes initial ranging for either of these reasons, it **MUST** use its Ranging SID instead of the initialization SID for the initial ranging process and not re-register. In the 3rd case, the CM **MUST** be able to transmit normally by 1,5 ms times the number of US channels in the TCS that have been changed within 1,5 ms of each other. For example, if the changes for 3 TDMA channels within the TCS take effect simultaneously, the CM would have 4,5 ms to make all of the changes. If the CM receives a data grant during this reconfiguration period, it **MAY** ignore the grant and re-request for the bandwidth.

Additionally, using the Ranging Required parameter in the new UCD message, the CMTS can force the CM to perform ranging prior to making any other transmissions using the parameters in the new UCD message. In certain cases, channel wide parameter changes (in particular, Modulation Rate or Center Frequency) may invalidate pre-equalization and synchronization parameters and normal operation may not be possible without re-ranging. If the CMTS changes the Modulation Rate or Center Frequency on an S-CDMA channel, it **MUST** force re-ranging using the Ranging Required parameter.

In the case of an S-CDMA channel, the first UCD message with a new Configuration Count and any subsequent UCD messages that may be sent prior to the first MAP with the new UCD Change Count **MUST** have an updated timestamp snapshot corresponding to the start time of that first MAP with the new UCD Change Count. Also on an S-CDMA channel the CMTS **MUST** maintain the continuity of the mini-slot and S-CDMA frame counters during the change in UCD parameters even if the size of a mini-slot is changed.

The CMTS **MUST NOT** transmit MAPs with the old UCD Change Count after transmitting the new UCD message.

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

It is possible for the change in upstream parameters to cause the upstream to change from a Type 1 upstream (see clause 10.2.3.6) to a Type 2, Type 3 or a Type 4 upstream. If this happens and the CM registered with a configuration file that enables 2.0 Mode (see clause 10.2.5.6), then the CM MUST operate in 2.0 Mode. If the upstream has changed to a Type 2 or Type 4 upstream, this means that any request the CM transmits in an opportunity in the MAP with the new Configuration Change Count or any subsequent MAP MUST be calculated by the CM in terms of IUCs 9 and 10, rather than IUCs 5 and 6. If the upstream has changed to a Type 2 or Type 4 upstream, the CMTS MUST issue grants using IUCs 9 and 10. However, if the upstream has changed to a Type 2 channel and the CM registered with a configuration file that disabled 2.0 Mode, then the CM MUST continue to calculate requests in terms of IUCs 5 and 6. If the CM registered with a configuration file that disabled 2.0 Mode, then the CMTS MUST issue grants using IUCs 5 and 6. If the CM registered with a configuration file that disables 2.0 Mode and the new parameters have changed the upstream to a Type 3 or Type 4 upstream, then the CM MUST immediately reinitialize the MAC layer and attempt registration. It should be understood by the network operator that changing a Type 1 upstream to a Type 3 or Type 4 upstream will cause a significant disruption of service for any CMs with 2.0 Mode disabled as well as any DOCSIS 1.x CMs that are using the channel. Also such CMs will only be able to resume operation if there is a Type 1 or Type 2 upstream available to them.

In Multiple Transmit Channel Mode, when implementing a UCD change on one channel, the CM MUST NOT impact upstream data transmission on other channels. In Multiple Transmit Channel Mode, the CM MUST remember requests that it has already made before the UCD change. For a CM operating in Multiple Transmit Channel Mode, the CMTS MUST remember the requests that the CM had already made. If not operating in Multiple Transmit Channel Mode, a CM discards knowledge of ungranted requests made prior to the UCD change. For a CM not operating in Multiple Transmit Channel Mode, a CMTS discards knowledge of ungranted requests received prior to the UCD change.

11.2 Dynamic Service Flow Changes

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 11-1.




Figure 11-1: 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 both the SFID and Transaction ID or SFID only. For a Dynamic Service Change that is modifying an Upstream Drop Classifier, the Service Flow is conceptually the NULL Service Flow and is not signaled in the message. A Transaction ID which is reused for other SFID(s) indicates that the other side terminated the previous transaction. If a Dynamic Service request message is received which refers to the same Transaction ID as one that has already been processed, but service flow(s) other than those locked in this transaction, the device MAY trigger a DSx Ended input to the state machine(s) of SF(s) involved in the previous transaction. 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. To help support transaction identifier uniqueness between 2 devices in different states, the CM or CMTS initiating the transaction SHOULD change the transaction identifier for each new initiated transaction. The CM or CMTS initiating the transaction MUST wait at least T10 to re-use the transaction identifier. The DSA/DSC transactions consist of a request/response/acknowledge sequence. In the case of a DSC message that is modifying an Upstream Drop Classifier, the acknowledge is not required and its absence does not result in a failed transaction. The DSD transactions consist of a request/response sequence. The response messages transmitted by the CM or CMTS MUST contain a confirmation code of okay unless some exception condition was detected. The acknowledge messages transmitted by the CM or CMTS MUST include the confirmation code in the response unless a new exception condition arises. A more detailed state diagram, including transition states, is shown in figure 11-2. The detailed actions for each transaction will be given in the following clauses.

11.2.1 Dynamic Service Flow State Transitions

The Dynamic Service Flow State Transition Diagram, figure 11-2, 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 signalling. 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 by the CM and CMTS 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 by the CM and CMTS 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 CM or CMTS MUST handle the event normally, by sending the SF Delete-Remote to the ongoing DSx Transaction and by initiating a DSD-Remote transaction. In addition, the CM or CMTS 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 CM or CMTS 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 is 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.

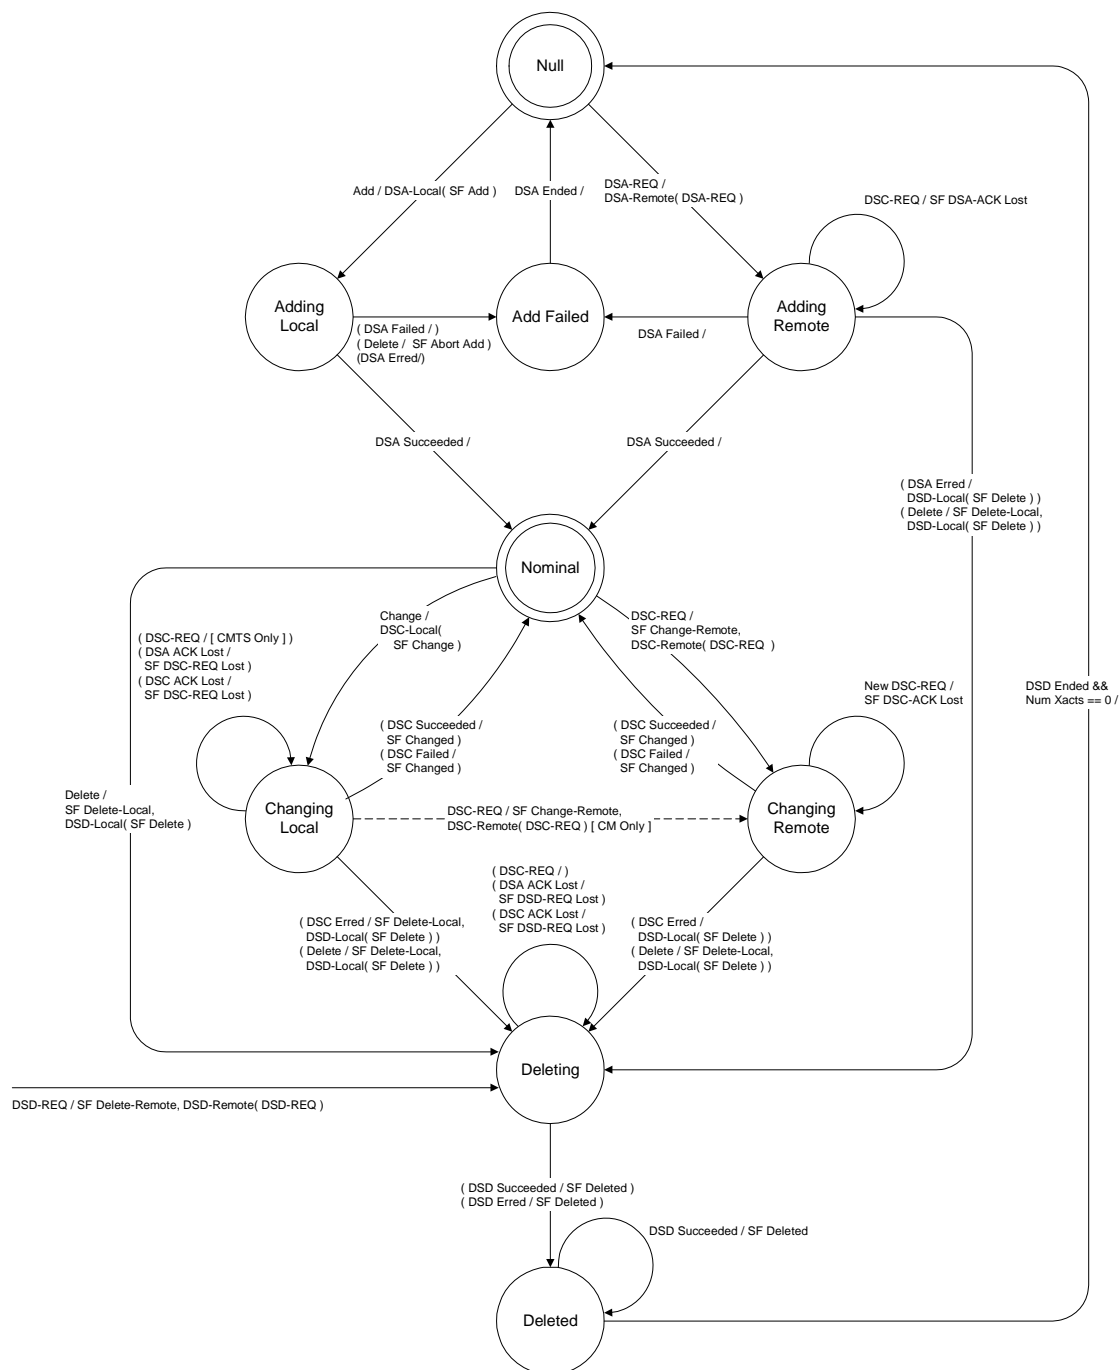


Figure 11-2: Dynamic Service Flow State Transition Diagram

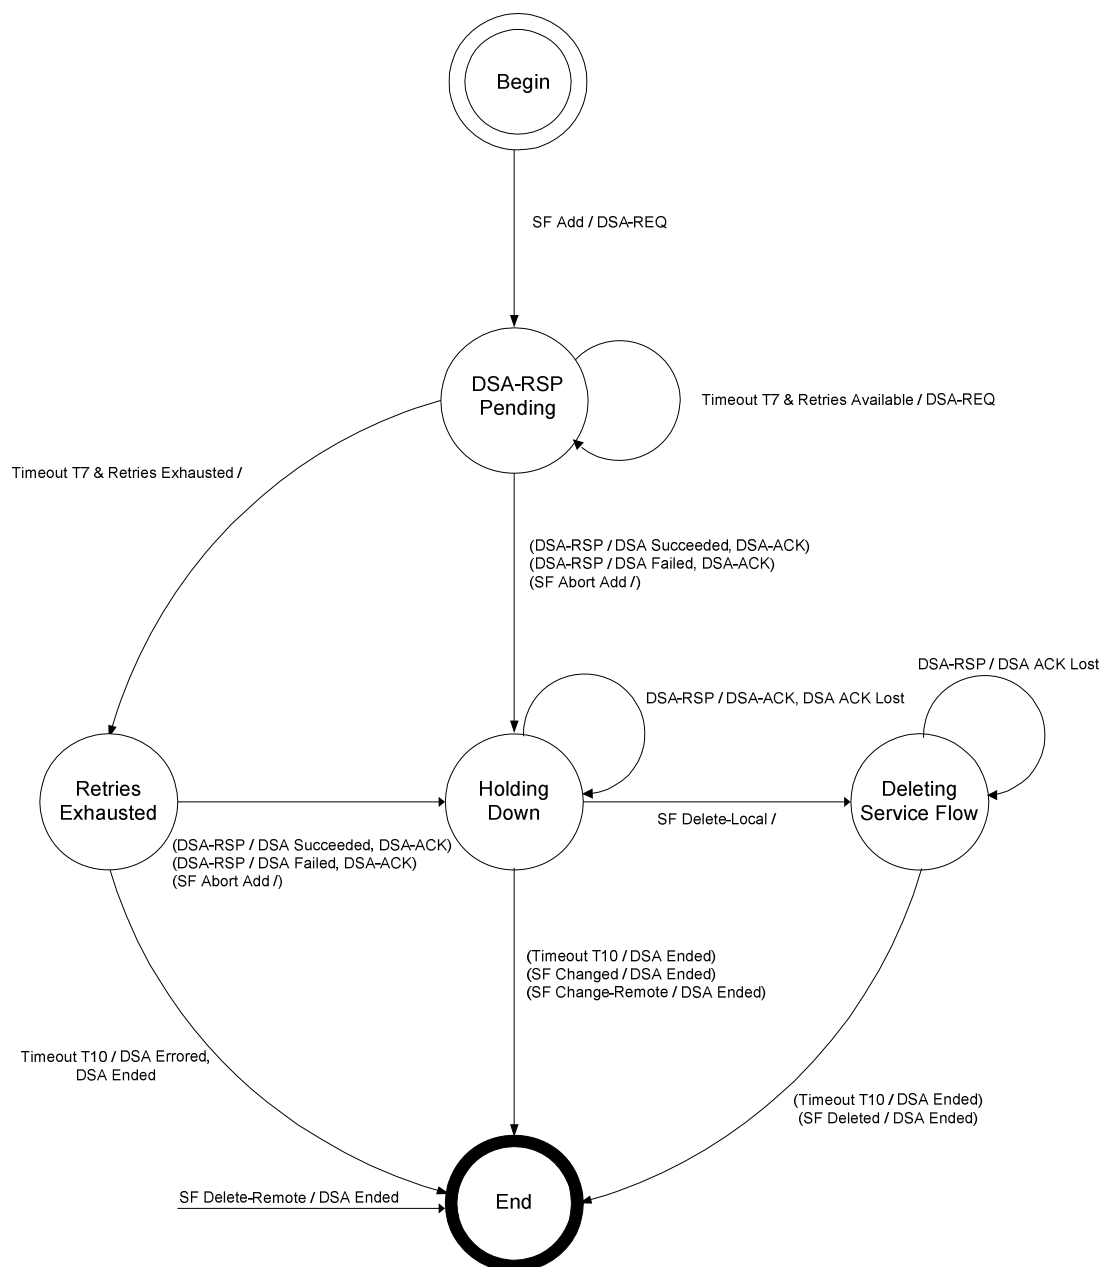


Figure 11-3: DSA-Locally Initiated Transaction State Transition Diagram

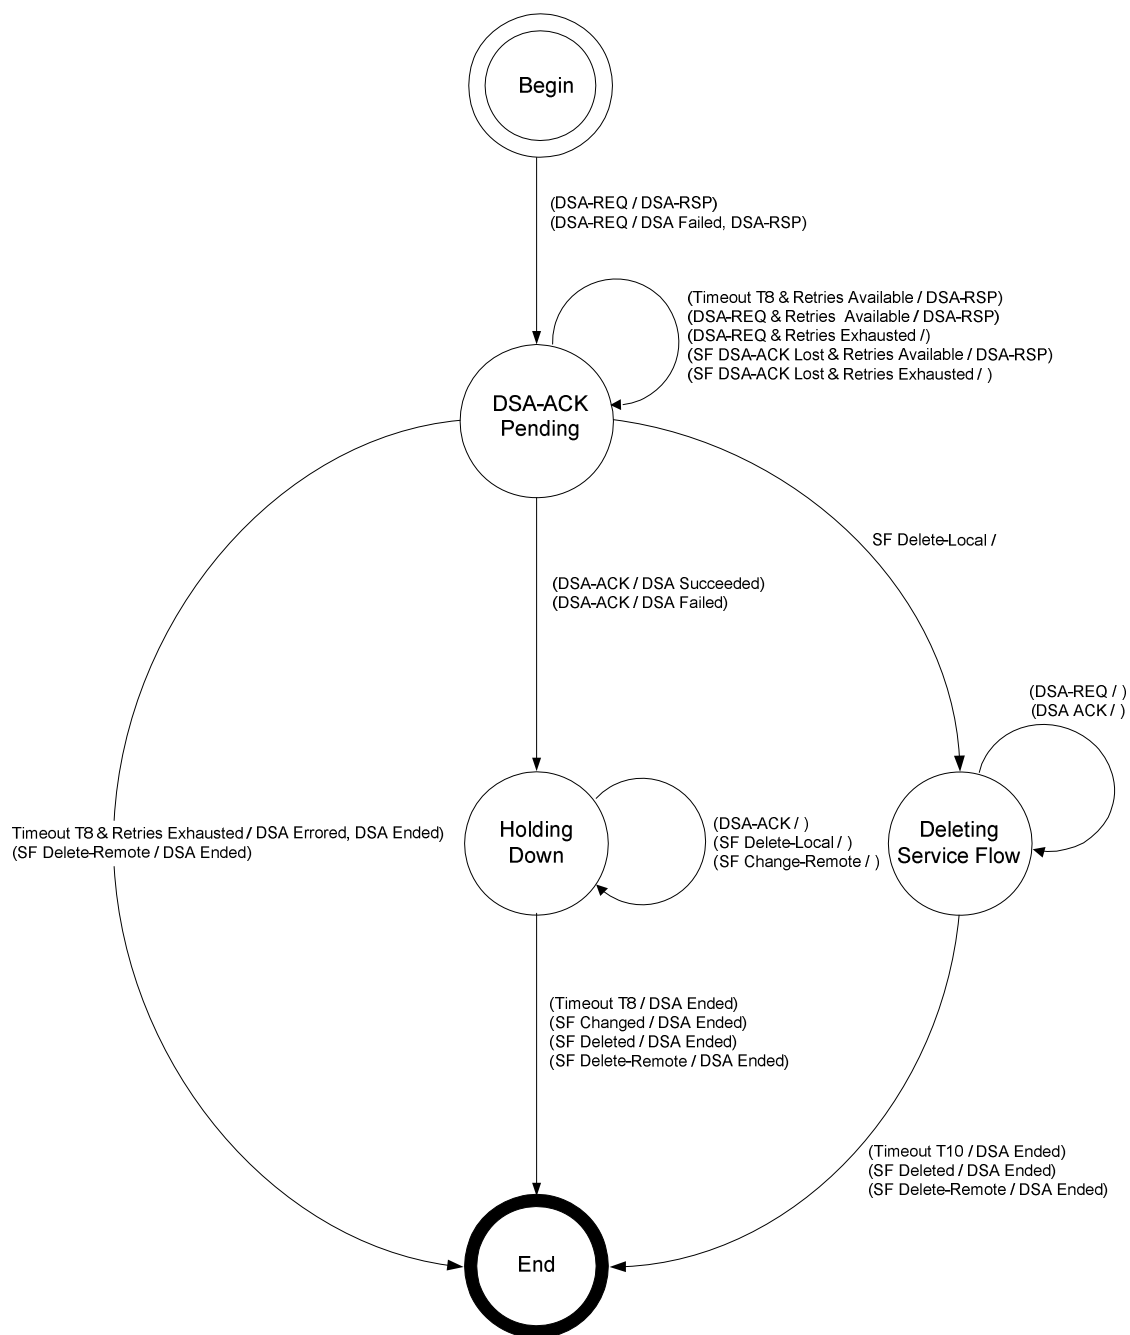


Figure 11-4: DSA-Remotely Initiated Transaction State Transition Diagram

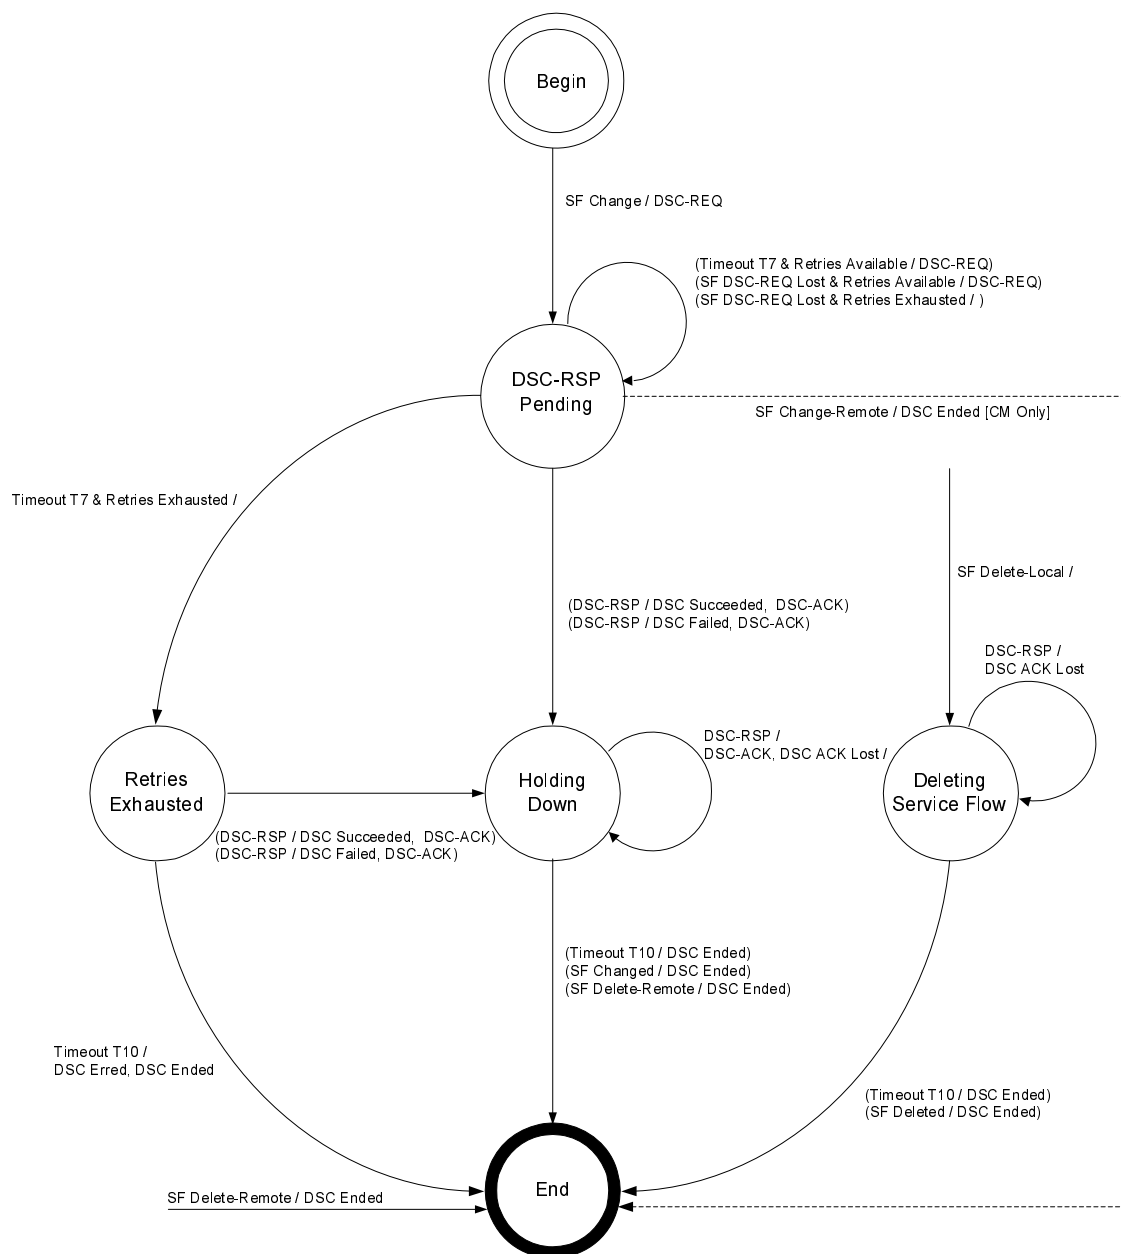


Figure 11-5: DSC-Locally Initiated Transaction State Transition Diagram

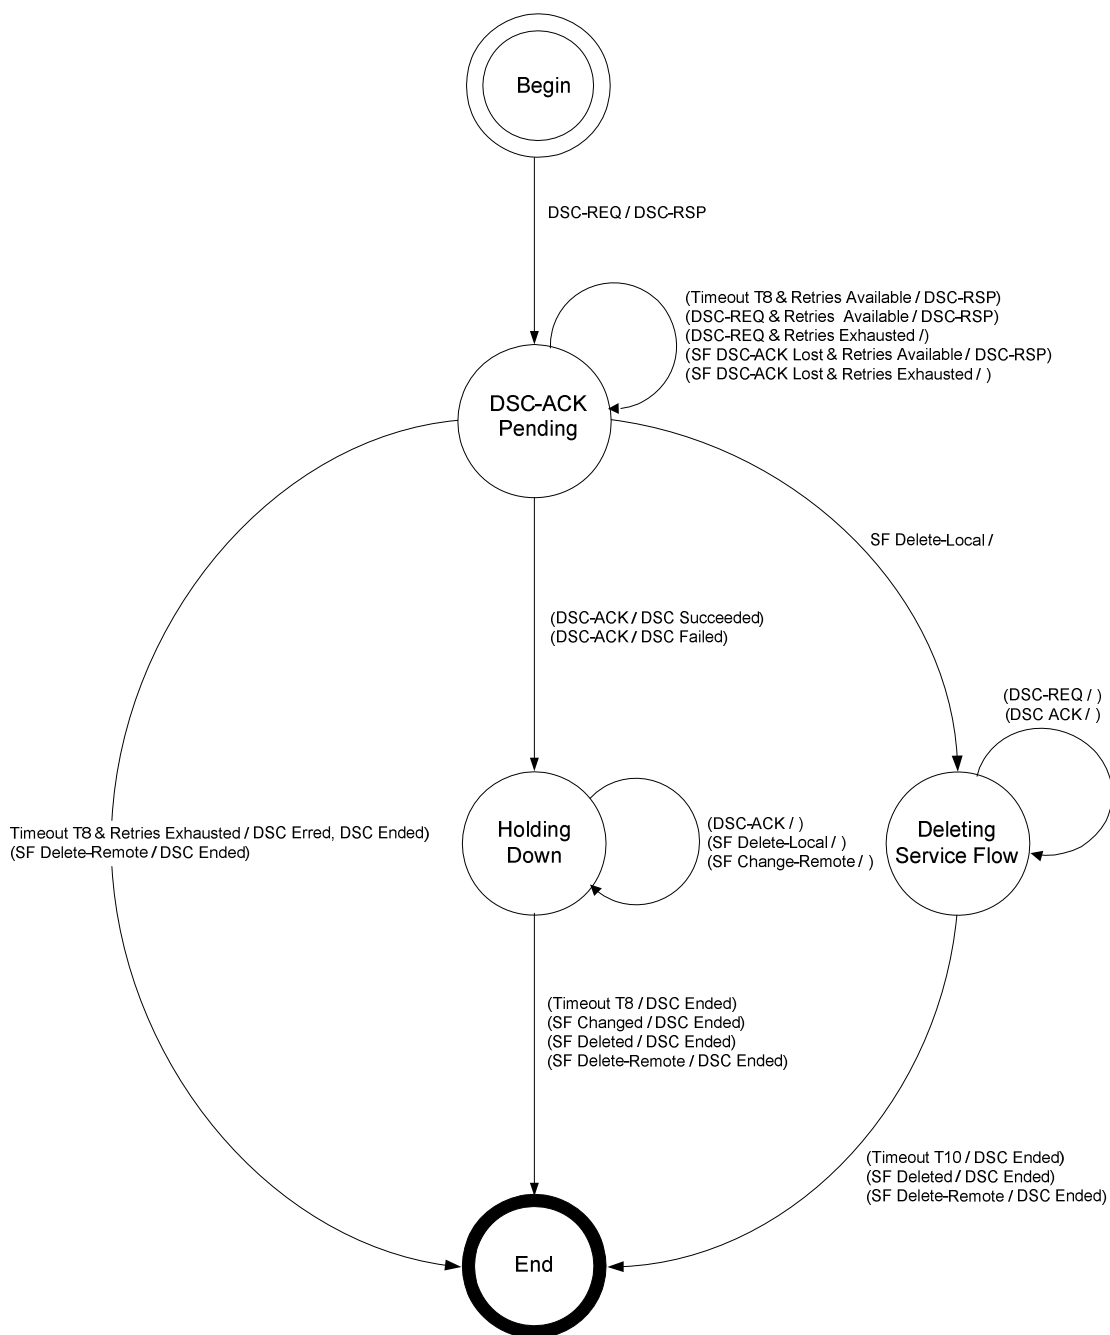


Figure 11-6: DSC-Remotely Initiated Transaction State Transition Diagram

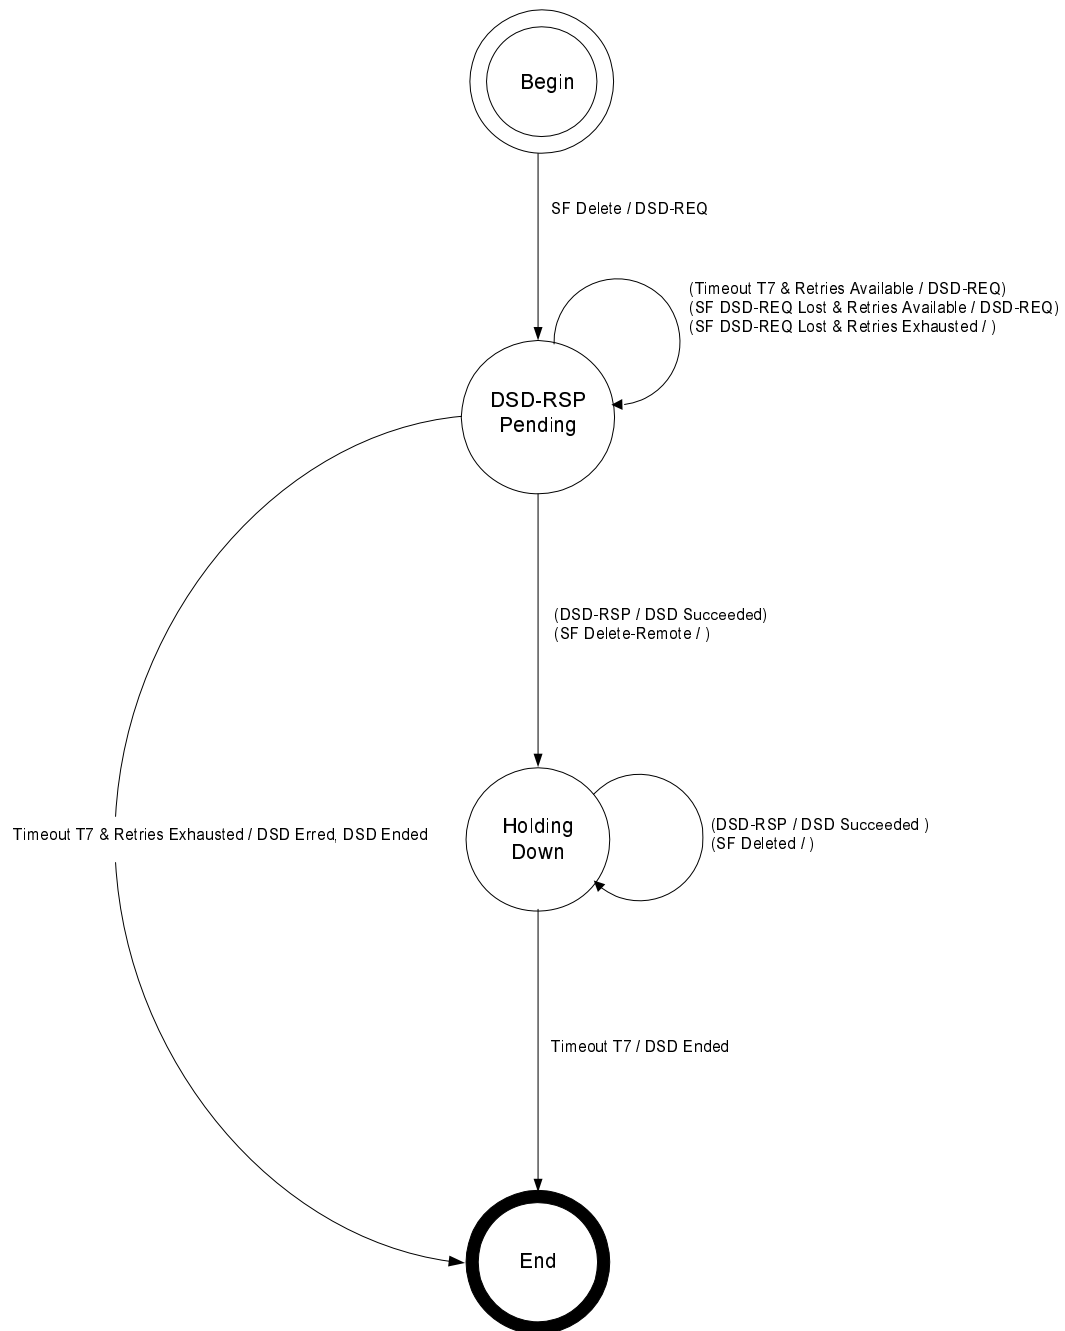


Figure 11-7: DSD-Locally Initiated Transaction State Transition Diagram

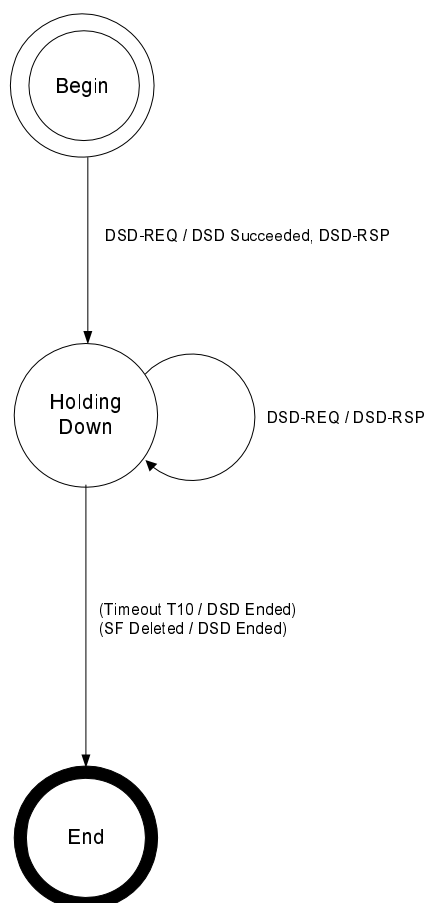


Figure 11-8: Dynamic Deletion (DSD) - Remotely Initiated Transaction State Transition Diagram

11.2.2 Dynamic Service Addition

11.2.2.1 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.

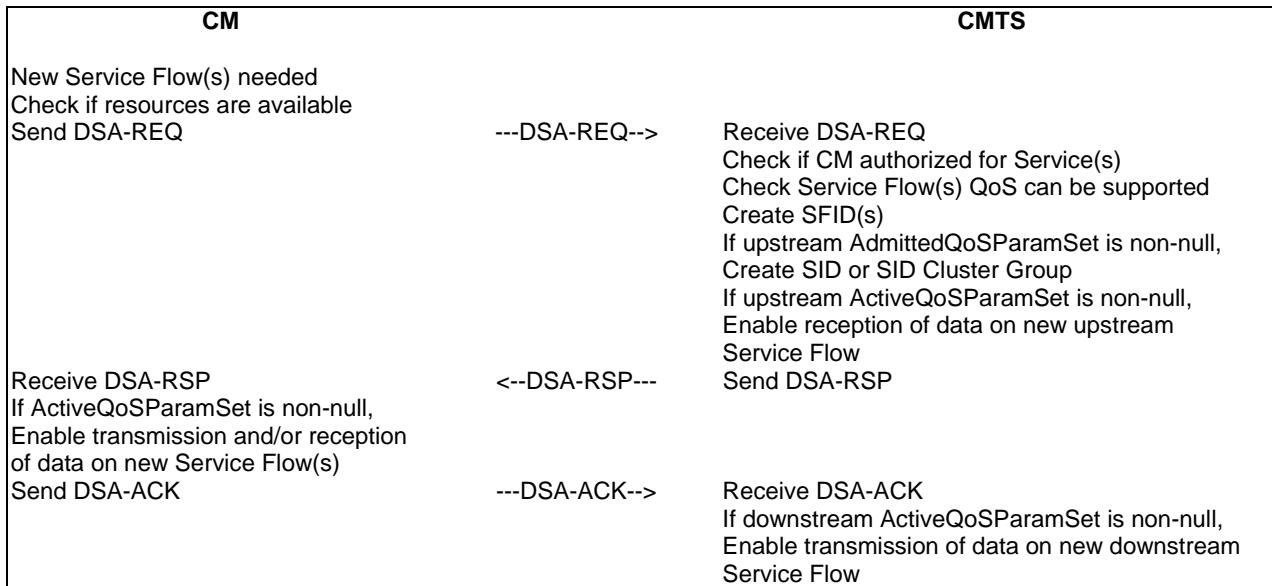


Figure 11-9: Dynamic Service Addition Initiated from CM

11.2.2.2 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). 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).

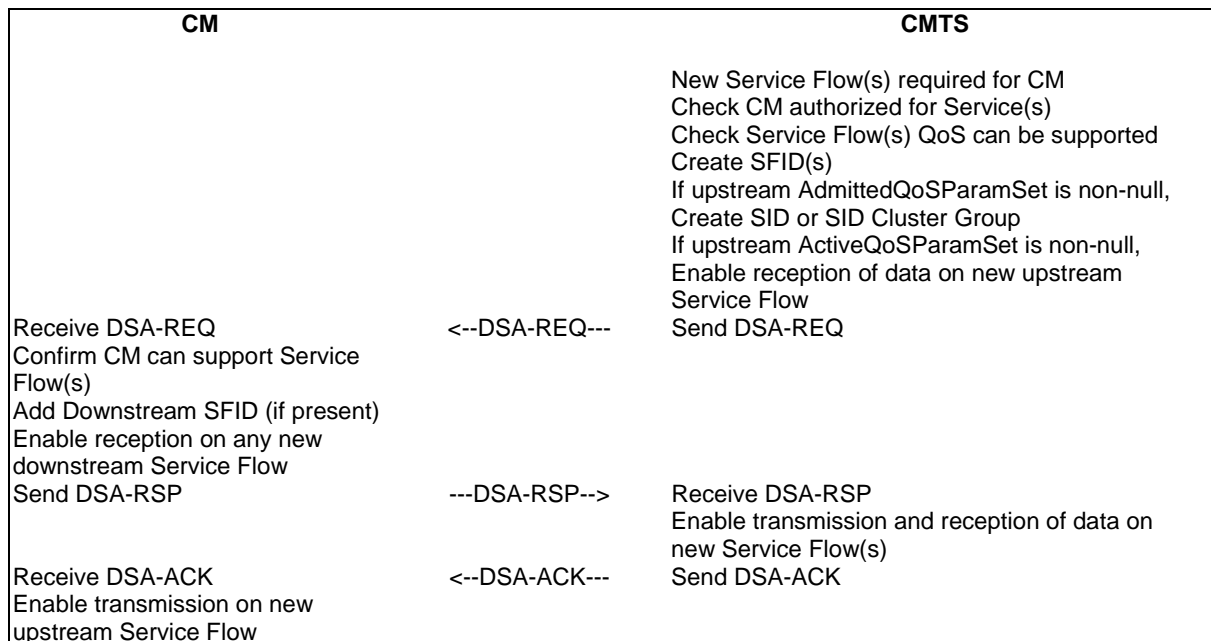


Figure 11-10: Dynamic Service Addition Initiated from CMTS

11.2.2.3 Dynamic Service Addition State Transition Diagrams




Figure 11-11: DSA-Locally Initiated Transaction Begin State Flow Diagram

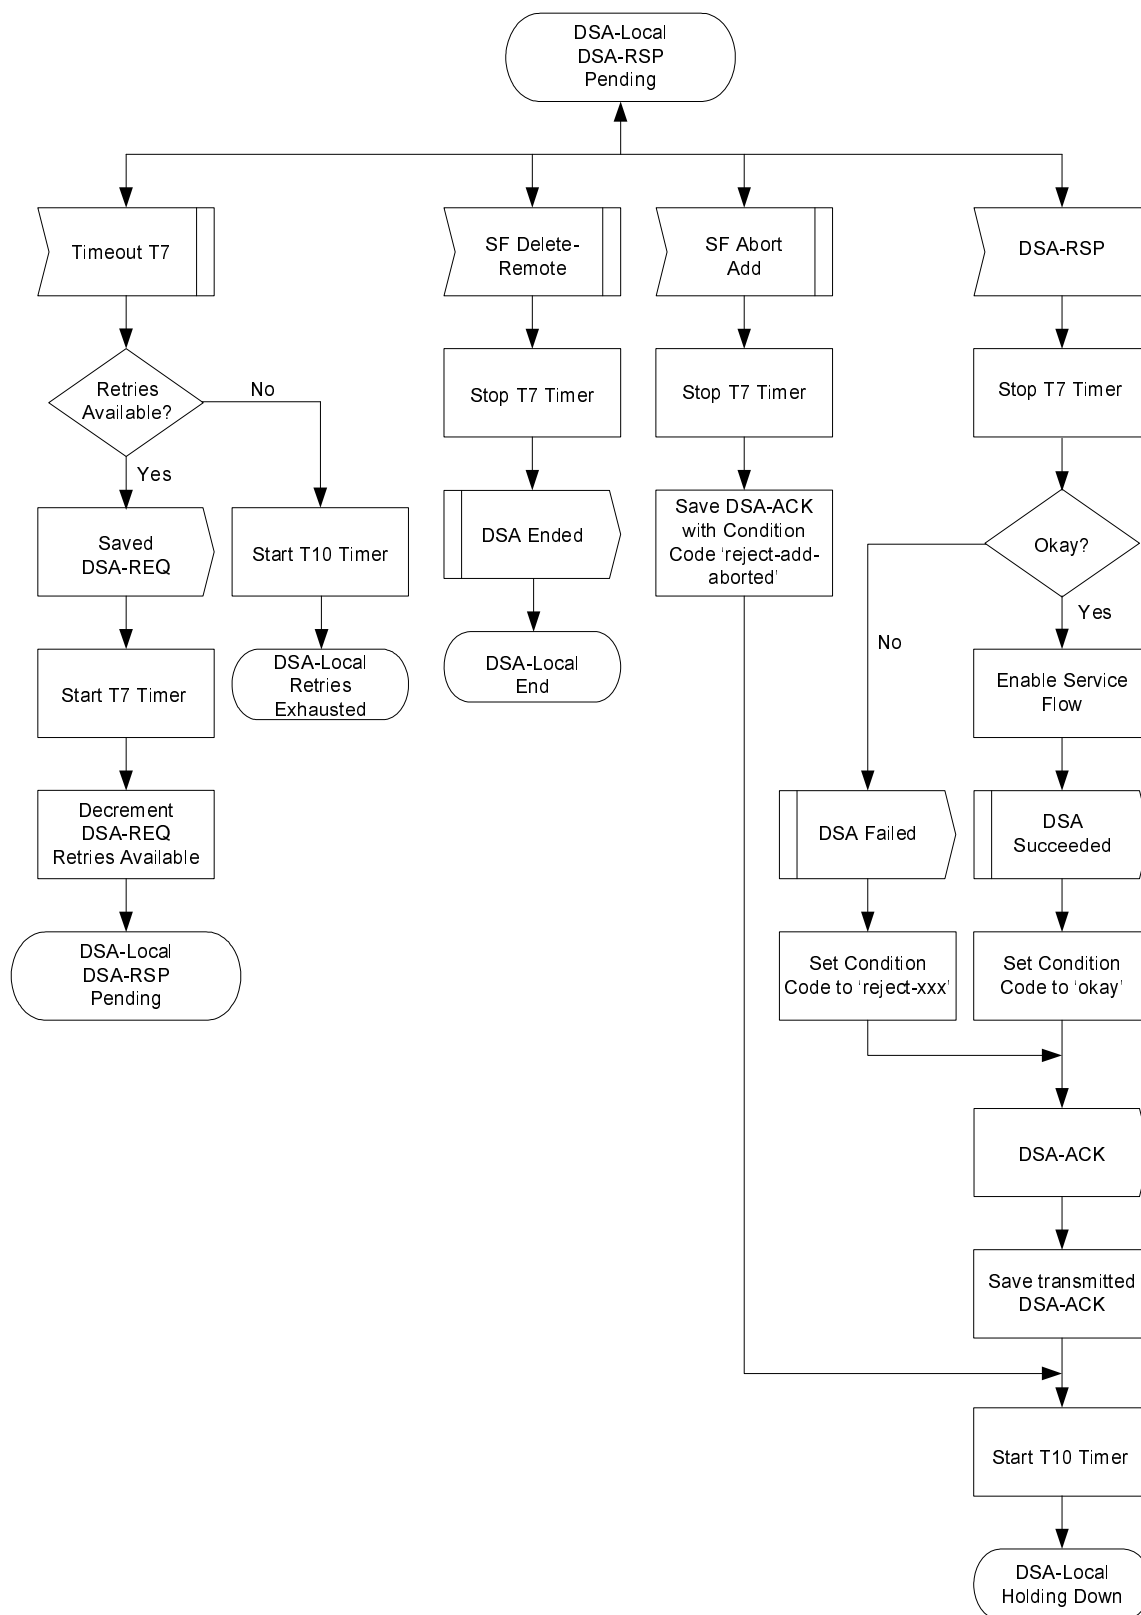


Figure 11-12: DSA-Locally Initiated Transaction DSA-RSP Pending State Flow Diagram

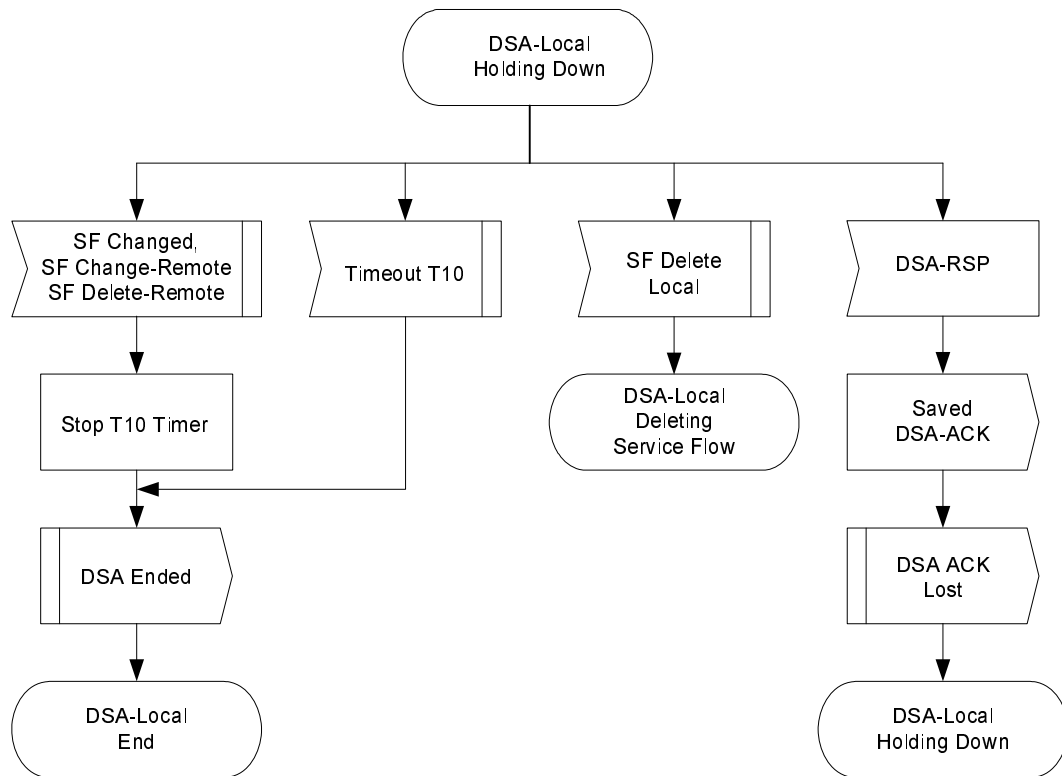


Figure 11-13: DSA-Locally Initiated Transaction Holding State Flow Diagram

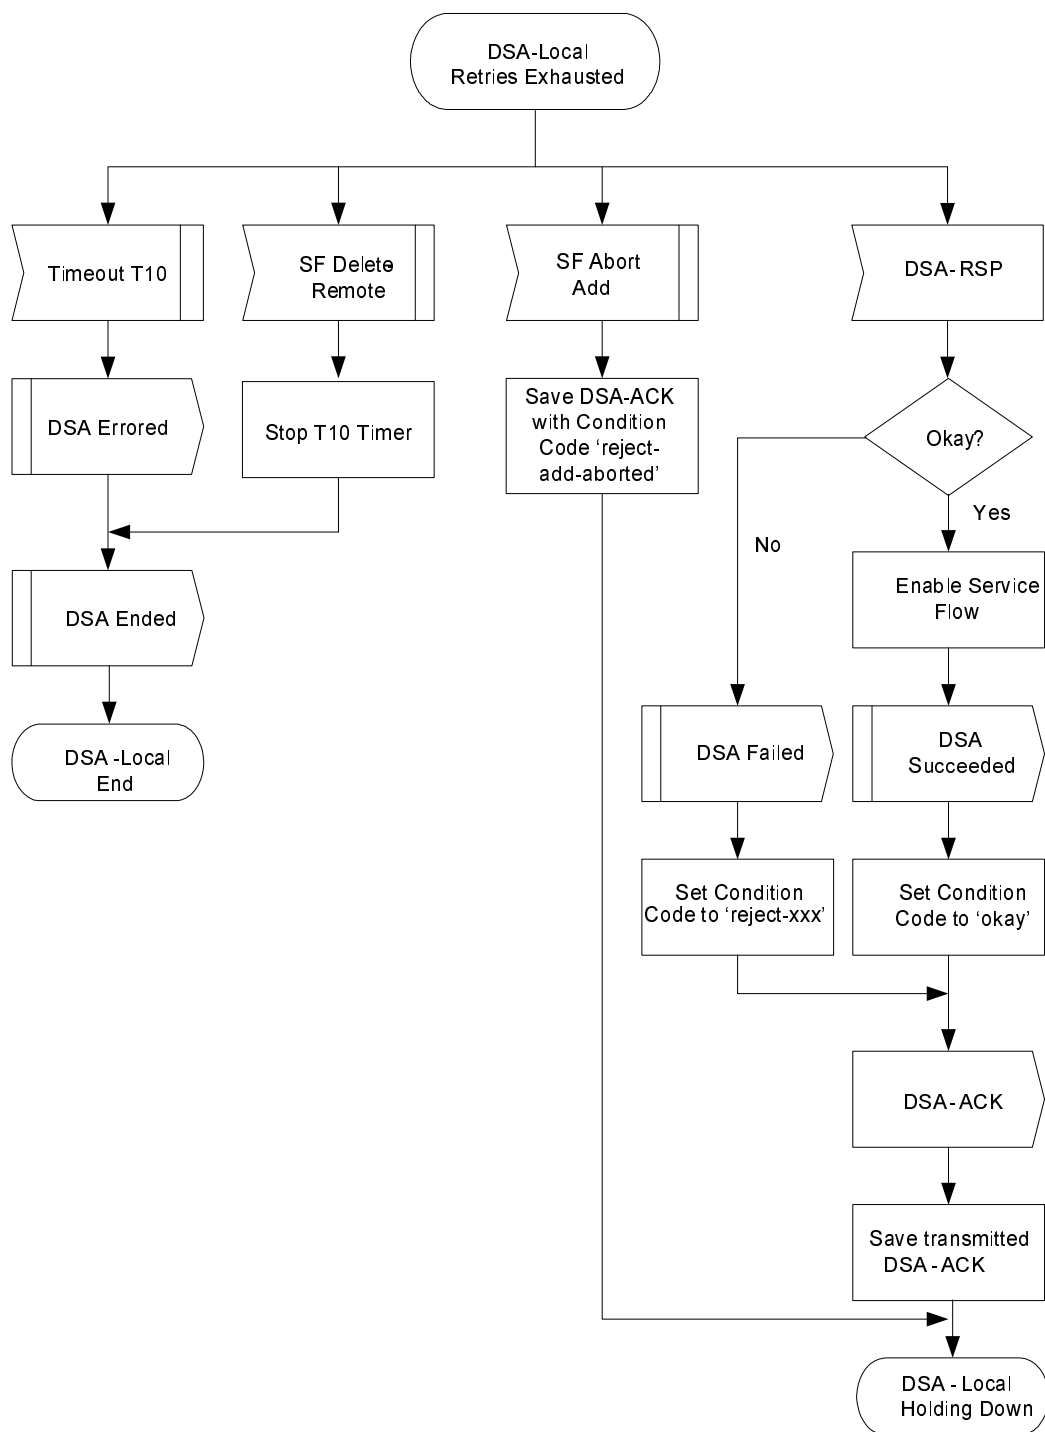


Figure 11-14: DSA-Locally Initiated Transaction Retries Exhausted State Flow Diagram

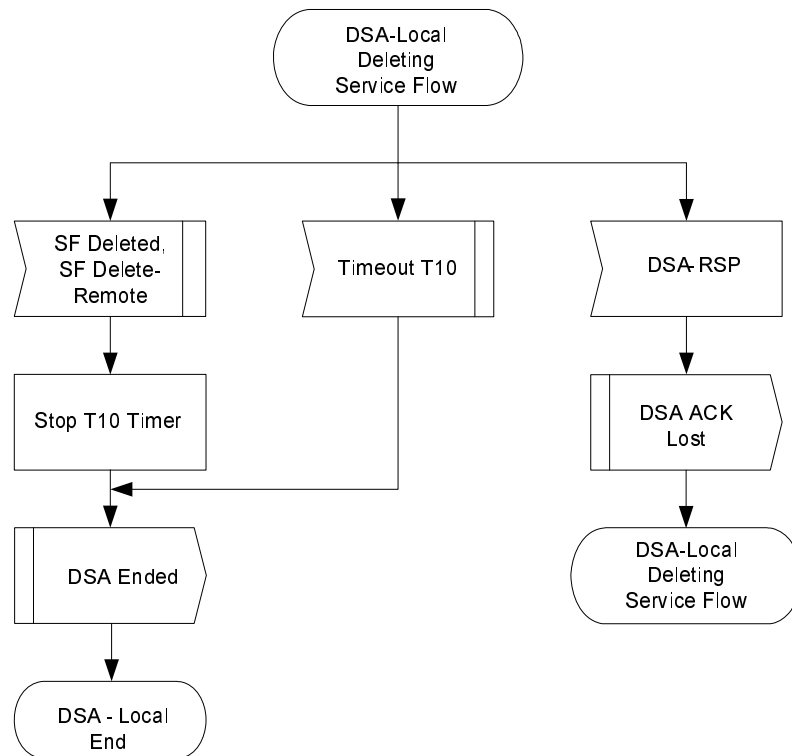


Figure 11-15: DSA-Locally Initiated Transaction Deleting Service Flow State Flow Diagram

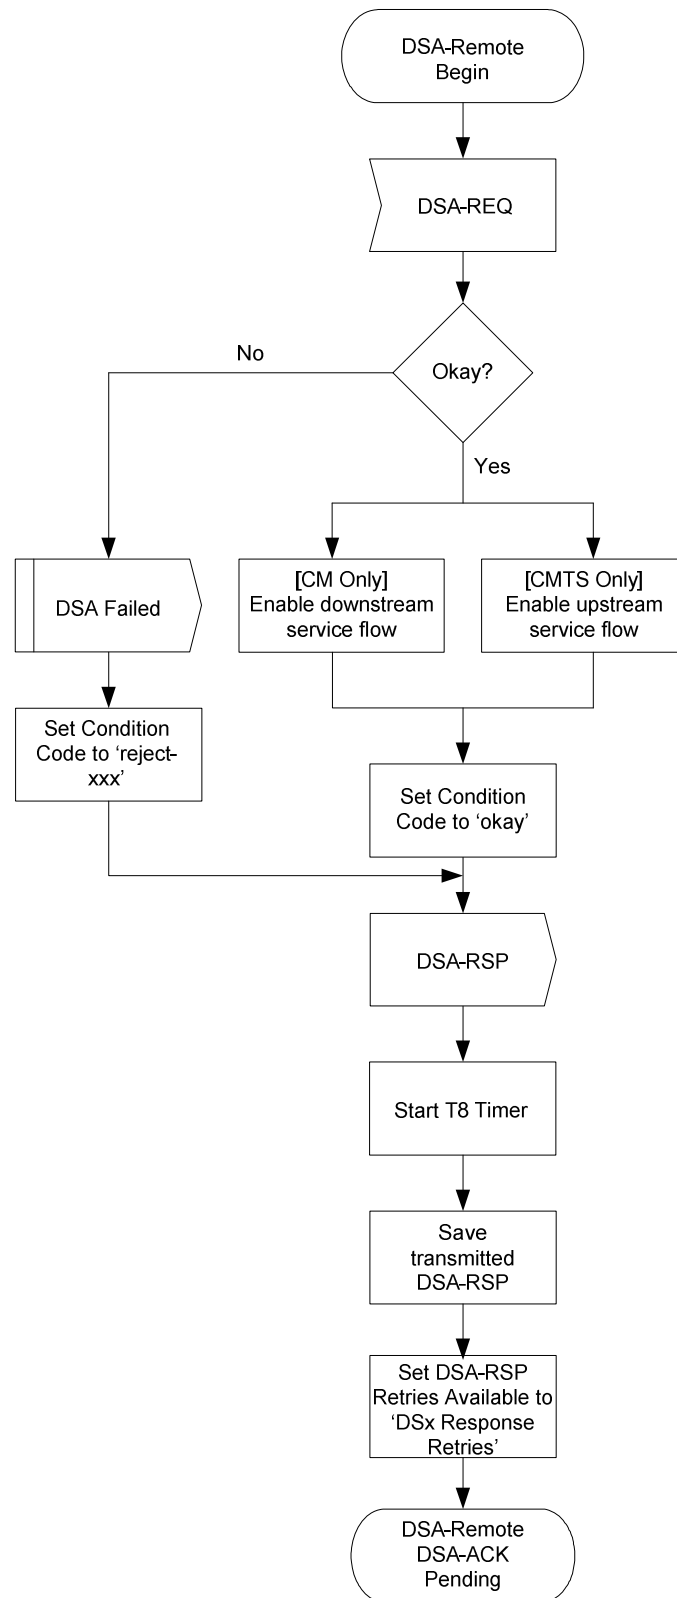


Figure 11-16: DSA-Remotely Initiated Transaction Begin State Flow Diagram

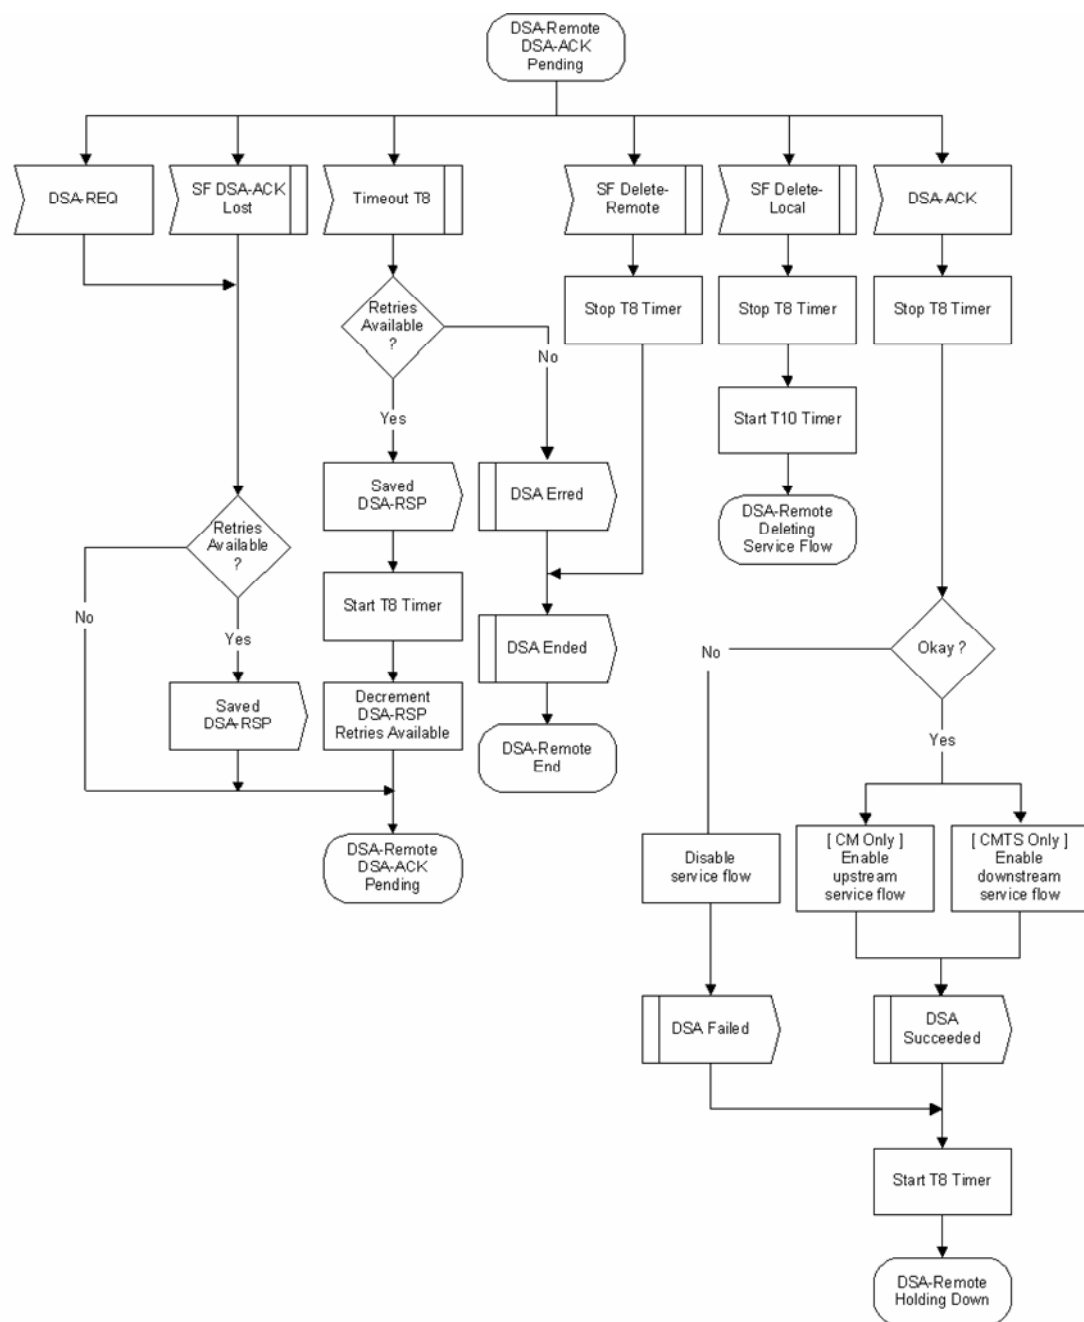


Figure 11-17: DSA-Remotely Initiated Transaction DSA-ACK Pending State Flow Diagram

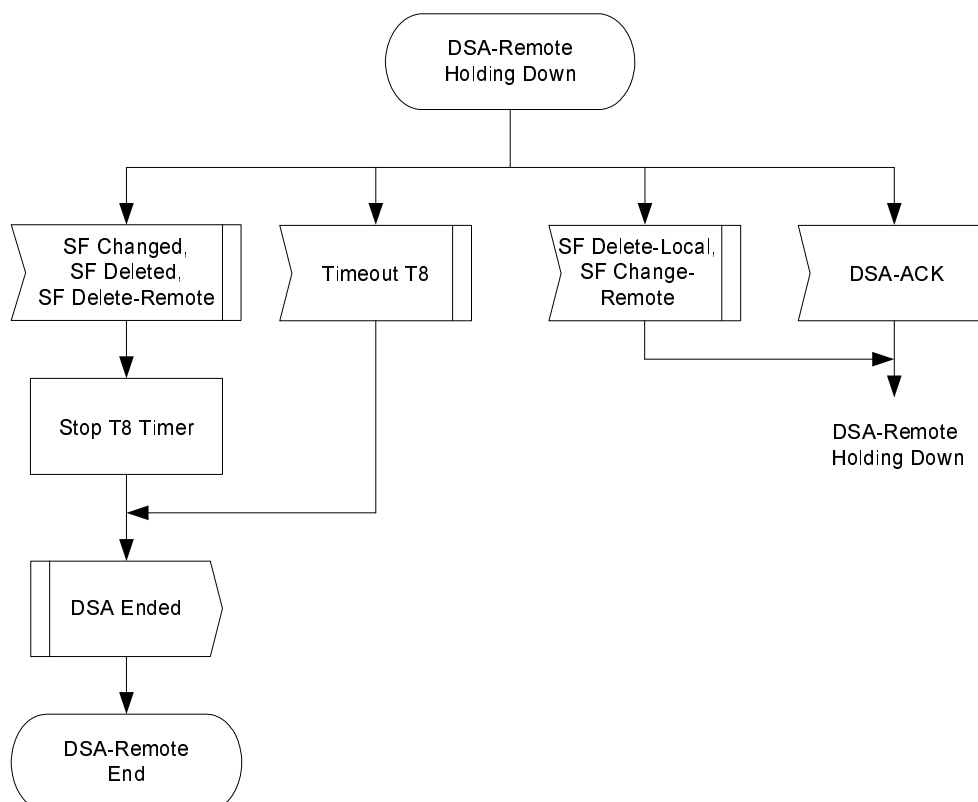


Figure 11-18: DSA-Remotely Initiated Transaction Holding Down State Flow Diagram

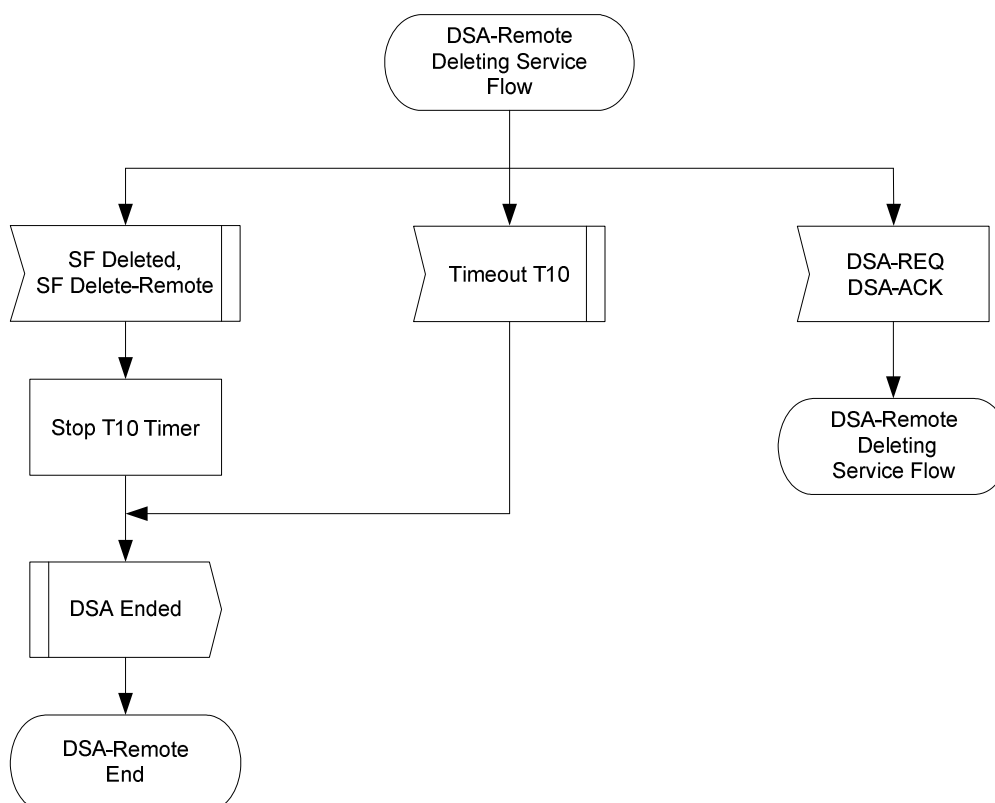


Figure 11-19: DSA-Remotely Initiated Transaction Deleting Service State Flow Diagram

11.2.3 Dynamic Service Change

The Dynamic Service Change (DSC) set of messages is used to modify the flow parameters associated with a Service Flow or a set of Upstream Drop Classifiers. Conceptually, Upstream Drop Classifiers are associated with a NULL Service Flow that is not signaled in the messages. Specifically, DSC can:

- Modify the Service Flow Specification or a set of Upstream Drop Classifiers.
- Add, Delete or Replace a Flow Classifier or a set of Upstream Drop Classifiers.
- 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. A single DSC message can modify multiple Upstream Drop Classifiers. If a CMTS is sending a DSC message that is modifying Upstream Drop Classifiers, it **MUST NOT** modify downstream or upstream Service Flow parameters. If a DSC is changing an Upstream Drop Classifier, then the term Service Flow used below, refers to the conceptual NULL Service Flow.

To prevent packet loss, any change to the bandwidth parameters of a Service Flow needs to be coordinated between the application generating the data and the DSC that modifies the Service Flow. 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 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. The CM changes Service Flow parameters on receipt of a DSC-REQ (CMTS-initiated transaction) or DSC-RSP (CM-initiated transaction).

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.

11.2.3.1 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
Receive DSC-REQ Validate Request Modify Service Flow	<----- DSC-REQ ----->	Service Flow Requires Modifying Send DSC-REQ
Send DSC-RSP	----- DSC-RSP ----->	Receive DSC-RSP Modify Service Flow
Receive DSC-ACK	<----- DSC-ACK ----->	Send DSC-ACK

Figure 11-20: CM-Initiated DSC

11.2.3.2 CMTS-Initiated Dynamic Service Change

A CMTS initiated DSC transaction that is changing Upstream Drop Classifiers does not require the CMTS to send a DSC-ACK after receiving a DSC-RSP from the CM. This is different from a CMTS initiated DSC transaction that is modifying a Service Flow and results from the fact that the CM cannot send a DSD if the transaction fails. The following paragraphs describe the DSC Transactions for a CMTS initiated DSC that is modifying a Service Flow versus a CMTS initiated DSC transaction that is modifying an Upstream Drop Classifier.

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 Modify Service Flow	<----- DSC-RSP ----->	Send DSC-RSP
Send DSC-ACK	----- DSC-ACK ----->	Receive DSC-ACK

Figure 11-21: CMTS-Initiated DSC modifying a Service Flow

A CMTS that needs to change an Upstream Drop Classifier performs the following operations.

The CMTS informs the CM of the additions or modifications to the Upstream Drop Classifiers 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 updates any state information that it is maintaining concerning the Upstream Drop Classifiers that the CM is using. The CMTS **MAY** send a Dynamic Service Change Acknowledgment (DSC-ACK). The CM **MUST NOT** delete the Upstream Drop Classifiers in the case that it does not receive a DSC-ACK message after sending the DSC-RSP.




Figure 11-22: CMTS-Initiated DSC modifying an Upstream Drop Classifier

11.2.3.3 Dynamic Service Change State Transition Diagrams




Figure 11-23: DSC-Locally Initiated Transaction Begin State Flow Diagram

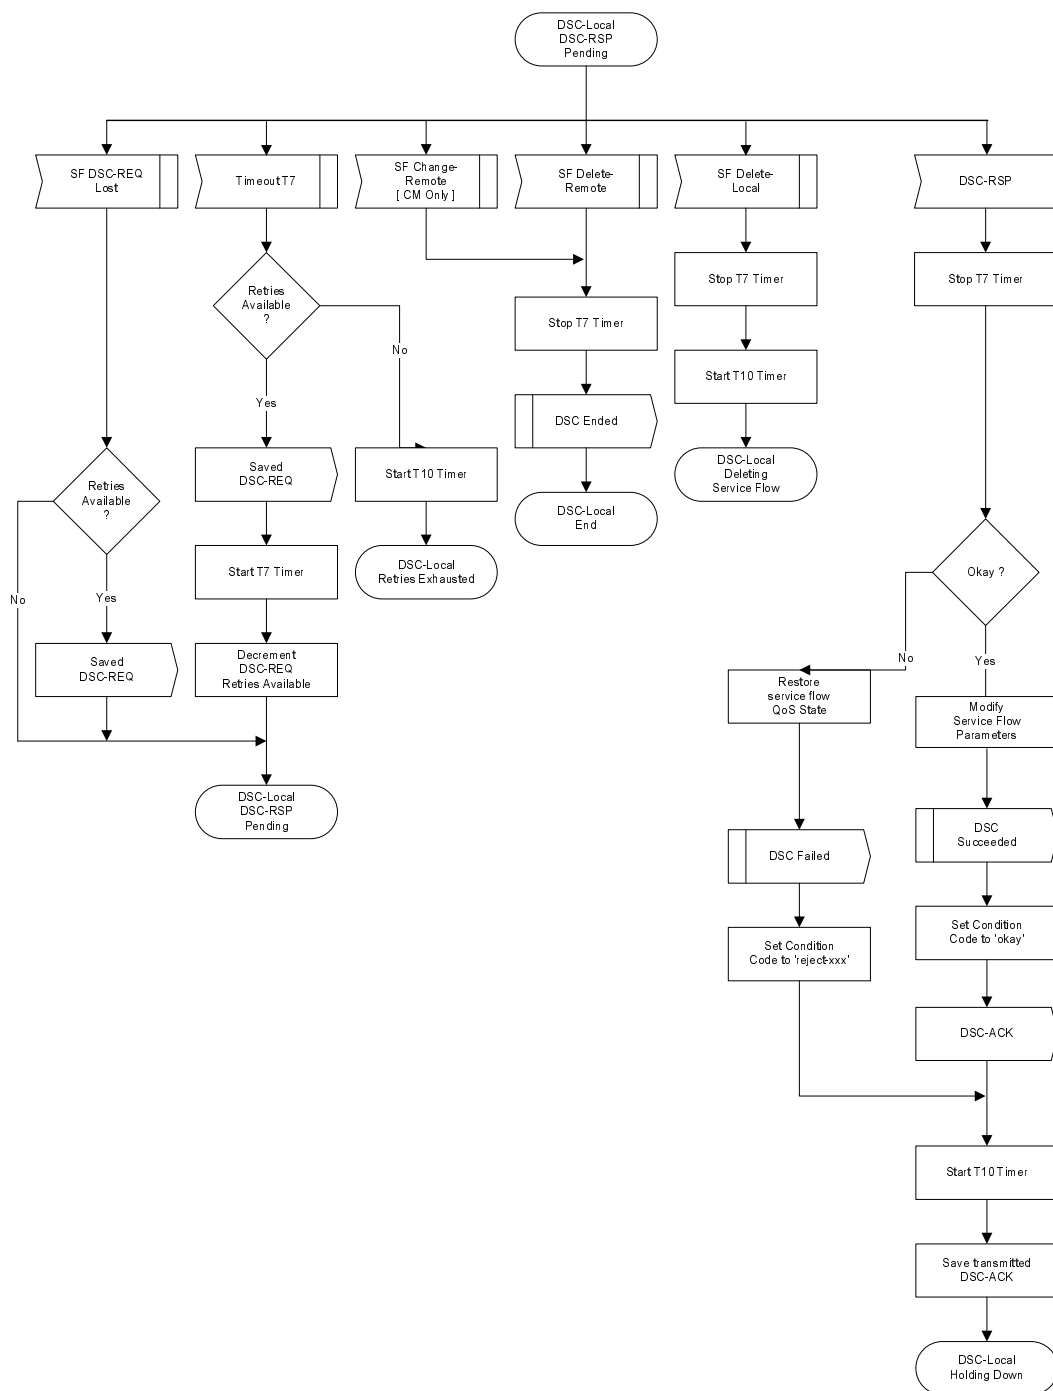


Figure 11-24: DSC-Locally Initiated Transaction DSC-RSP Pending State Flow Diagram

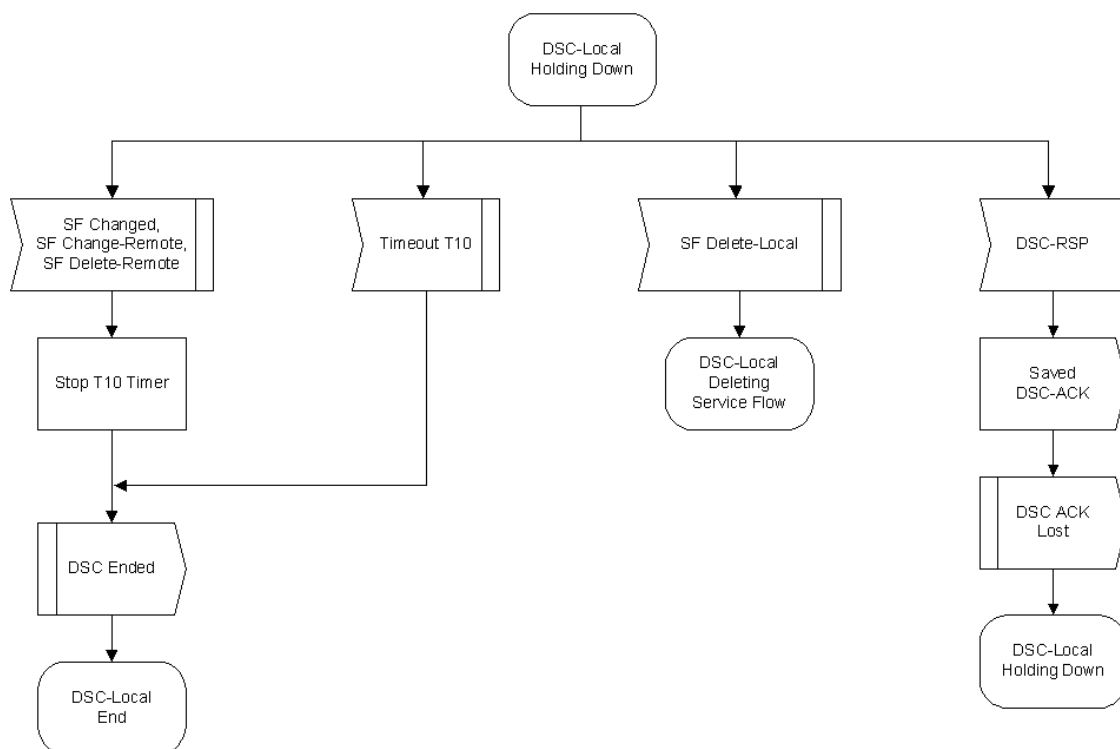


Figure 11-25: DSC-Locally Initiated Transaction Holding Down State Flow Diagram

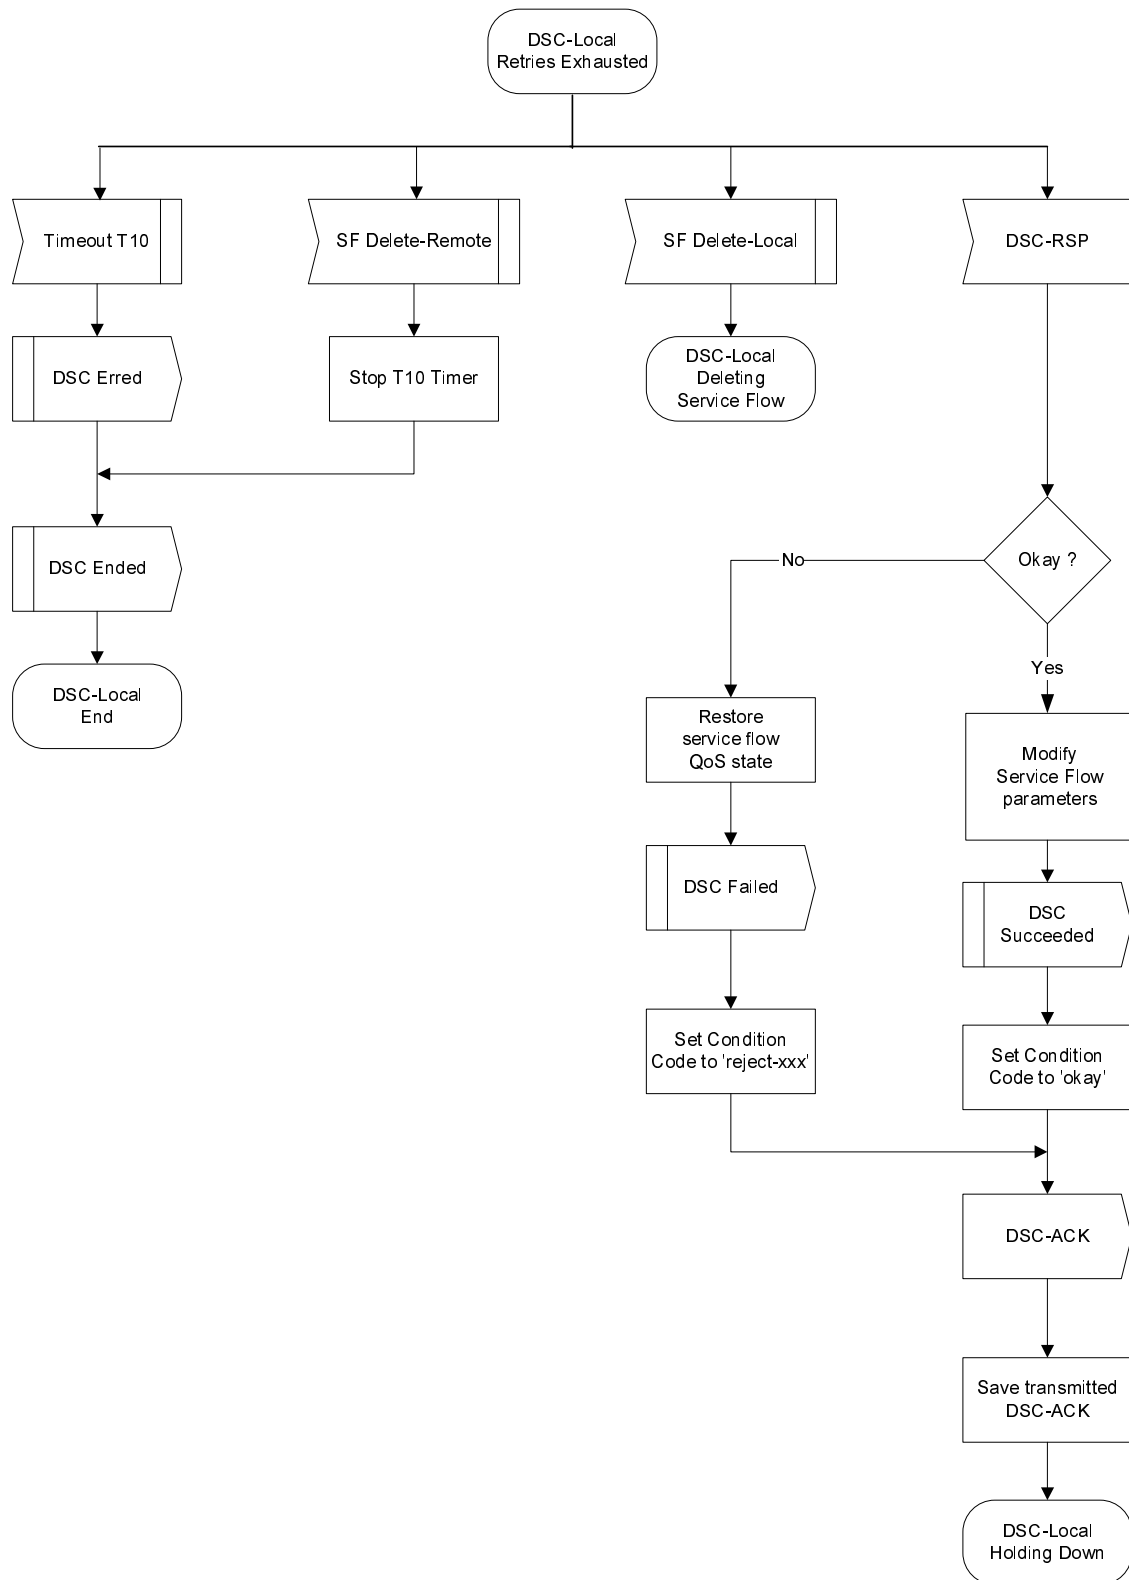


Figure 11-26: DSC-Locally Initiated Transaction Retries Exhausted State Flow Diagram

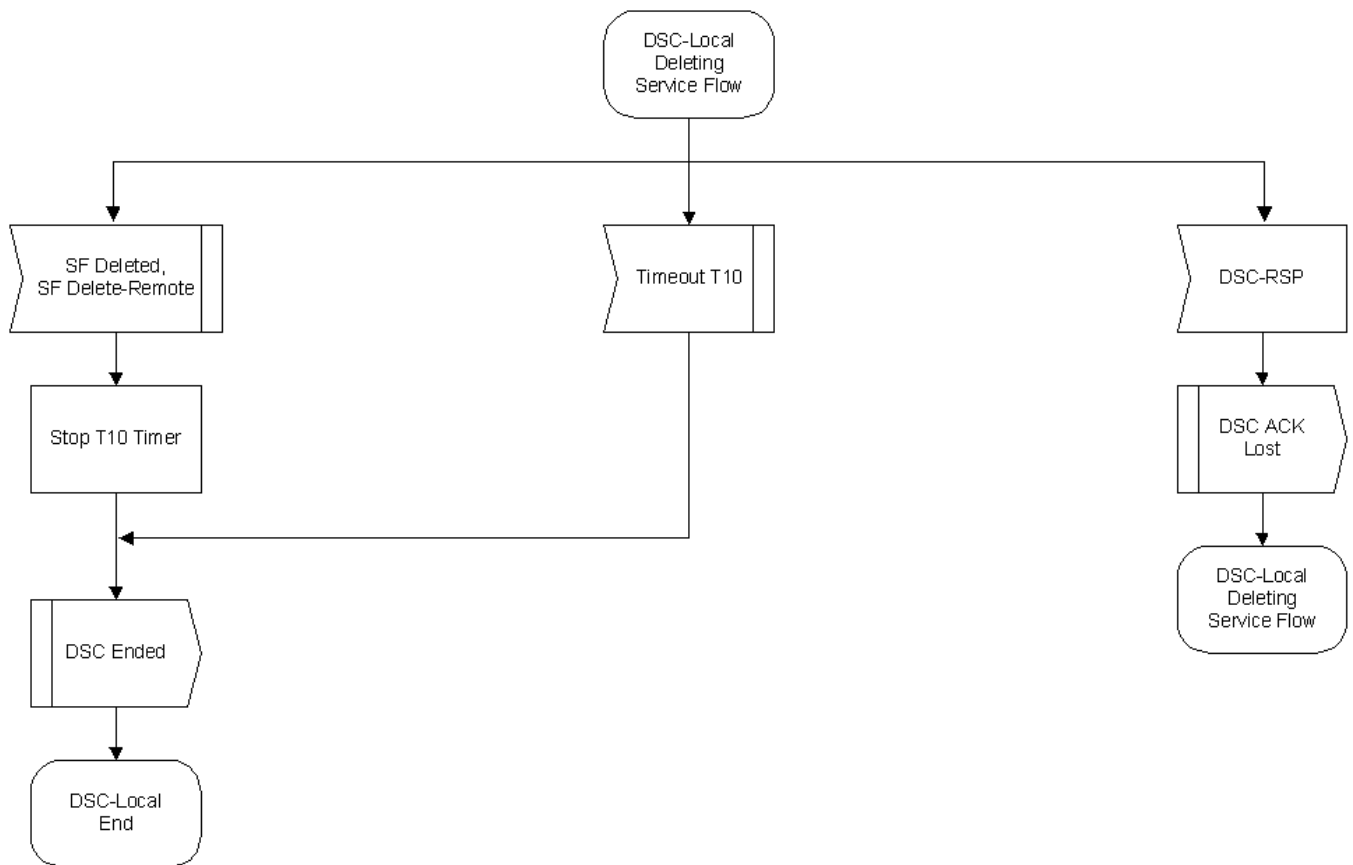


Figure 11-27: DSC-Locally Initiated Transaction Deleting Service Flow State Flow Diagram

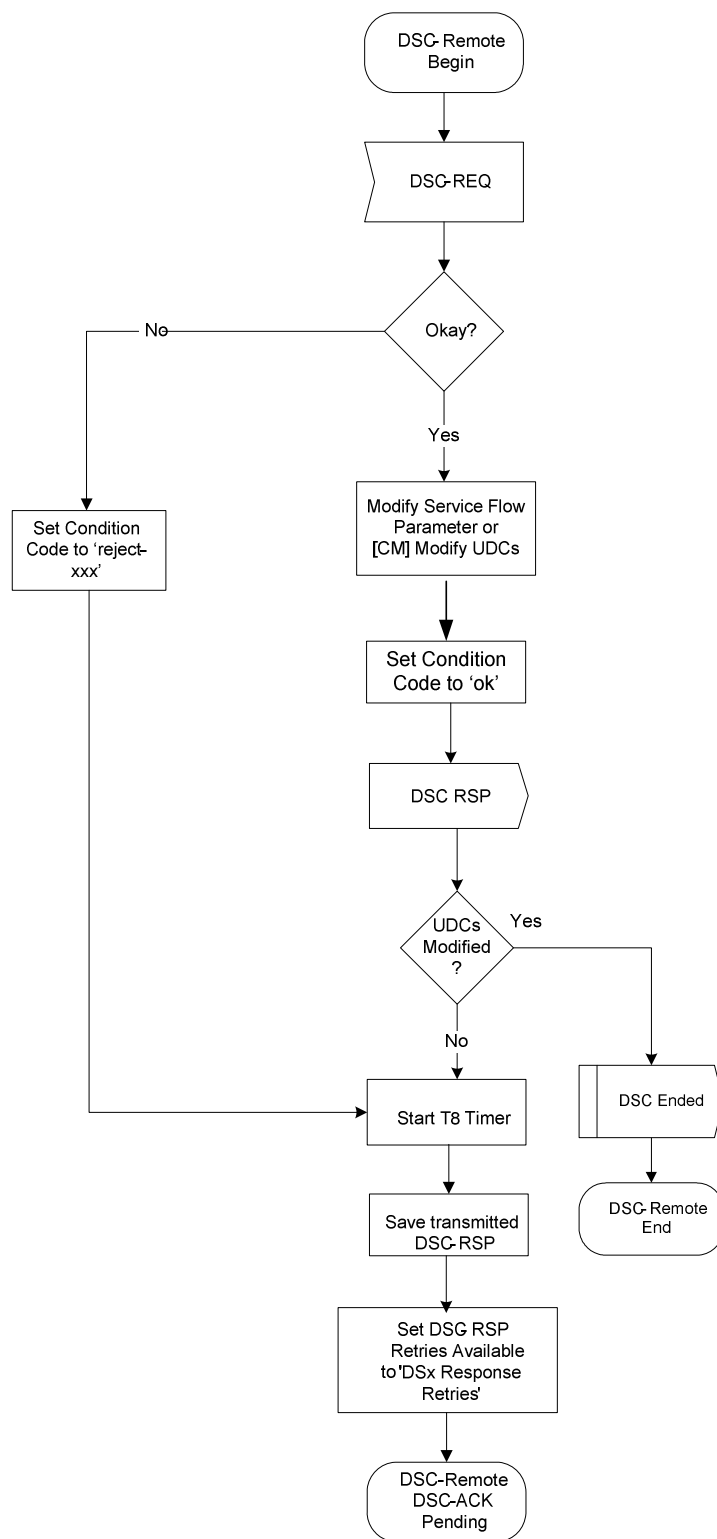


Figure 11-28: DSC-Remotely Initiated Transaction Begin State Flow Diagram

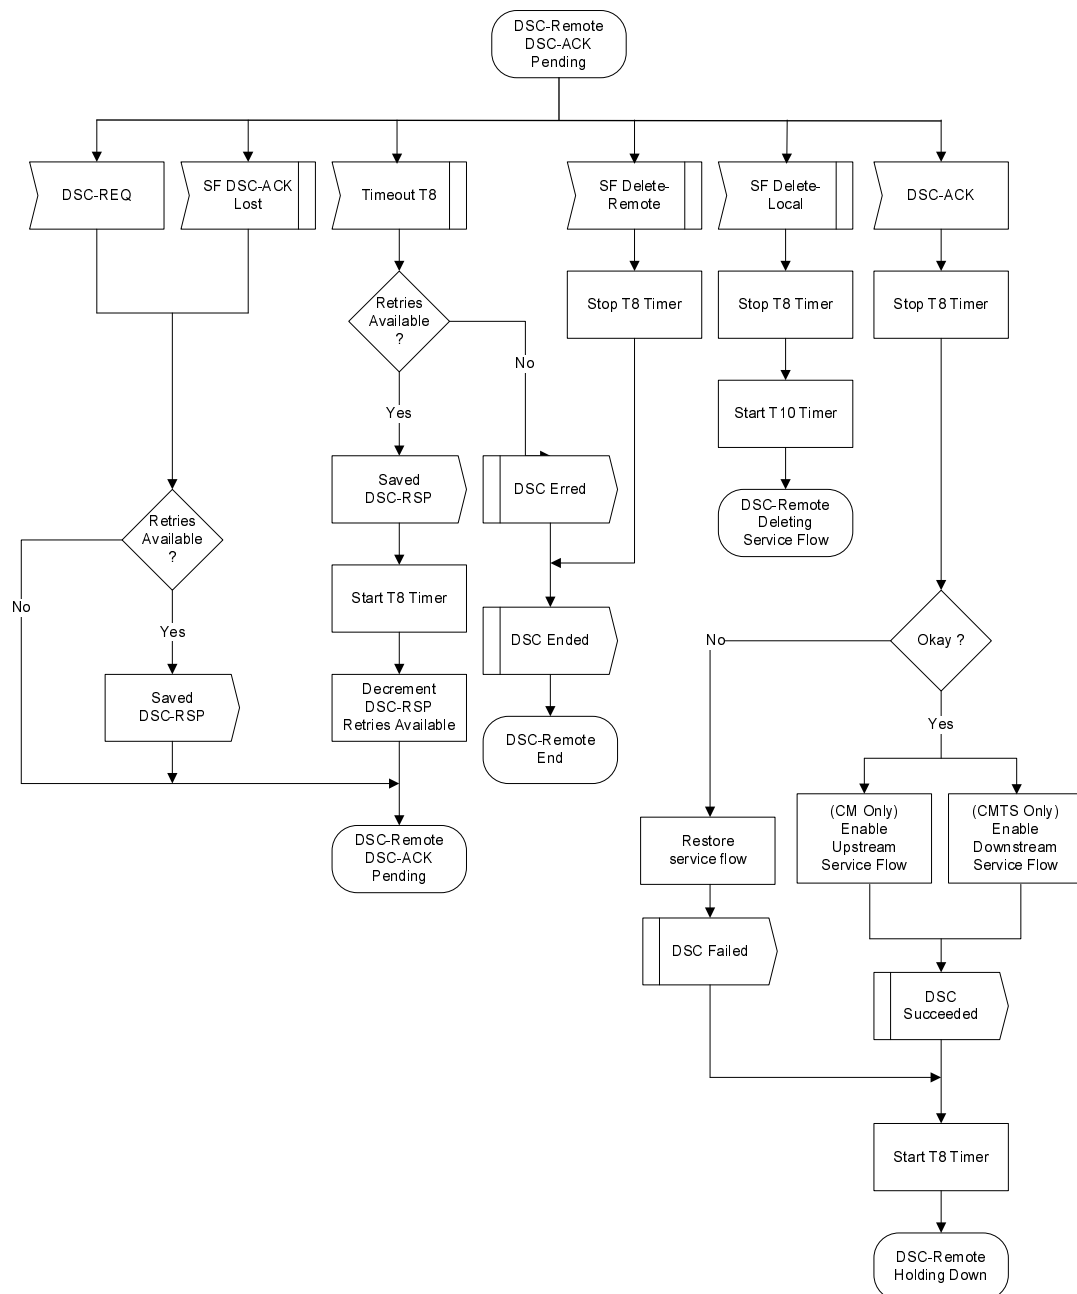


Figure 11-29: DSC-Remotely Initiated Transaction DSC-ACK Pending State Flow Diagram

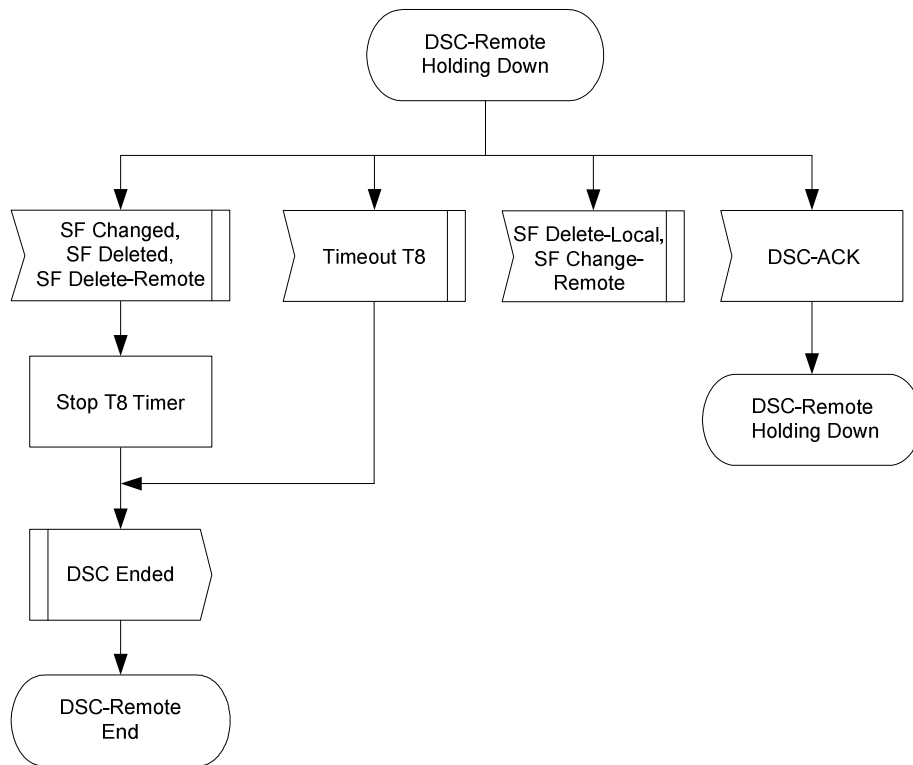


Figure 11-30: DSC-Remotely Initiated Transaction Holding Down State Flow Diagram

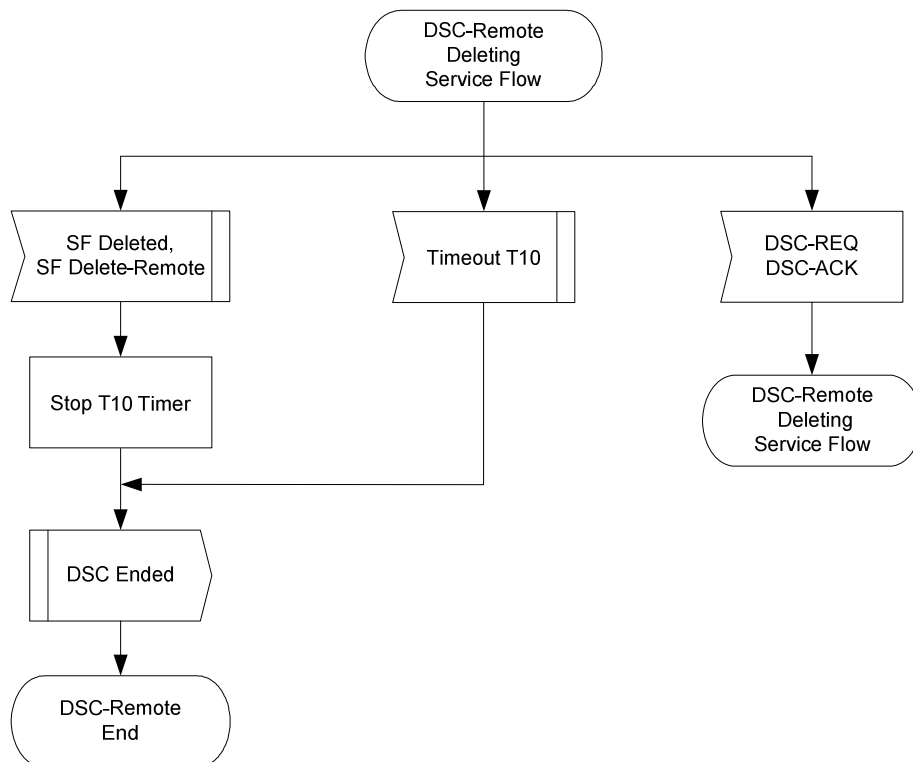


Figure 11-31: DSC-Remotely Initiated Transaction Deleting Service Flow State Flow Diagram

11.2.4 Dynamic Service Deletion

Any service flow can be deleted with the Dynamic Service Deletion (DSD) messages. When a Service Flow (either provisioned or dynamically created) is deleted, all resources associated with it are released, including classifiers, PHS Rules and SID Clusters. If a Primary Service Flow of a CM is deleted, that CM is de-registered and **MUST** re-register. However, the deletion of a provisioned Service Flow other than the Primary Service Flow **MUST NOT** cause a CM to re-register.

11.2.4.1 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) Send DSD-RSP
Receive DSD-RSP	<--DSD-RSP---	

Figure 11-32: Dynamic Service Deletion Initiated from CM

11.2.4.2 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) Send DSD-REQ
Receive DSD-REQ Delete Service Flow(s) Send DSD-RSP	<---DSD-REQ--- ---DSD-RSP-->	Receive DSD-RSP

Figure 11-33: Dynamic Service Deletion Initiated from CM

11.2.4.3 Dynamic Service Deletion State Transition Diagrams




Figure 11-34: DSD-Locally Initiated Transaction Begin State Flow Diagram

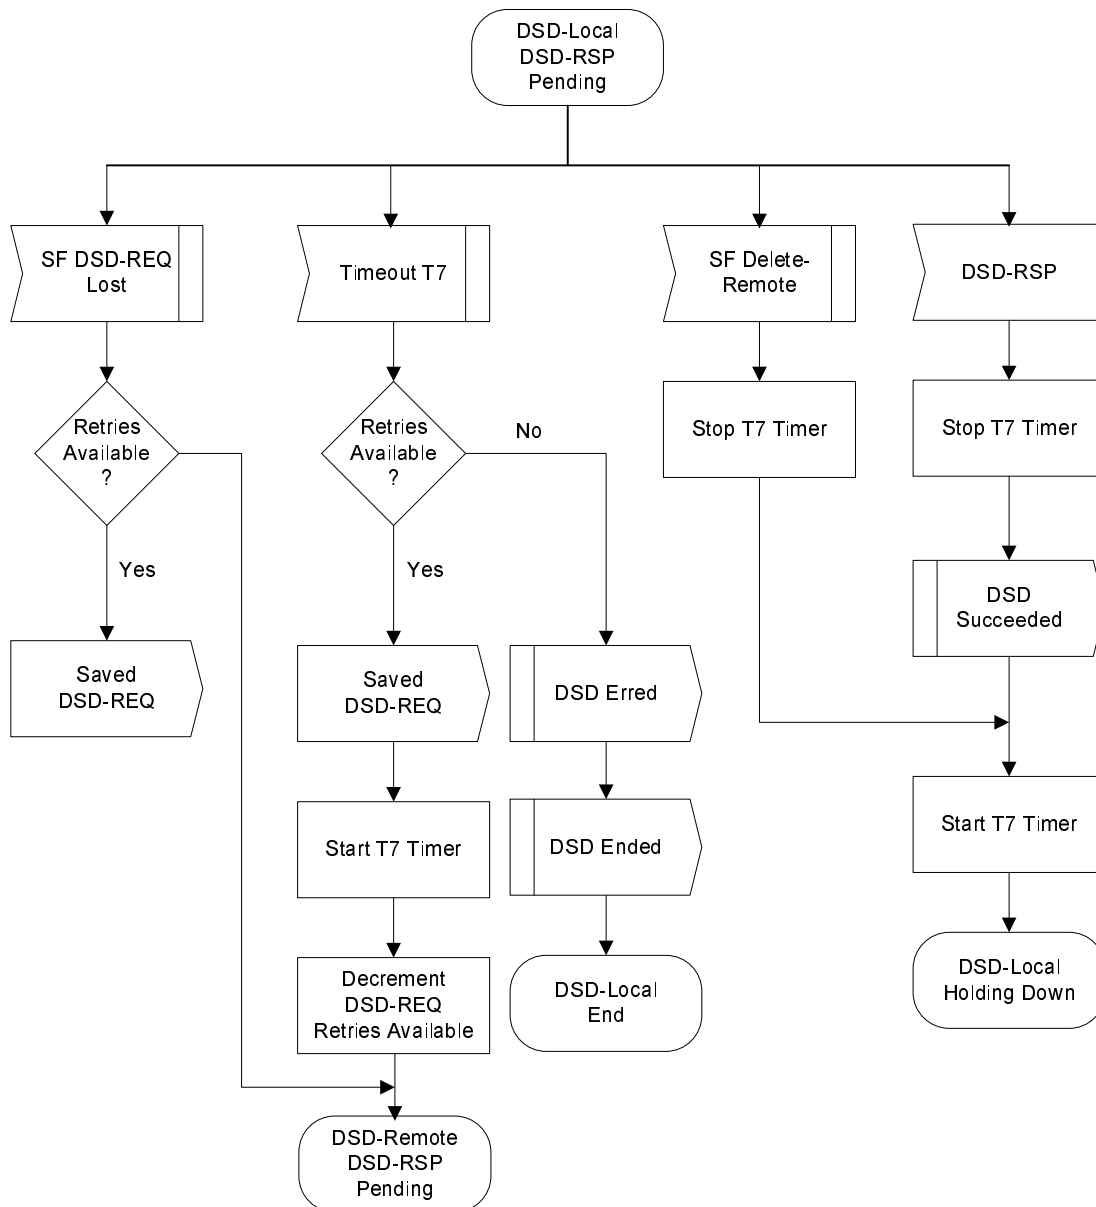


Figure 11-35: DSD-Locally Initiated Transaction DSD-RSP Pending State Flow Diagram

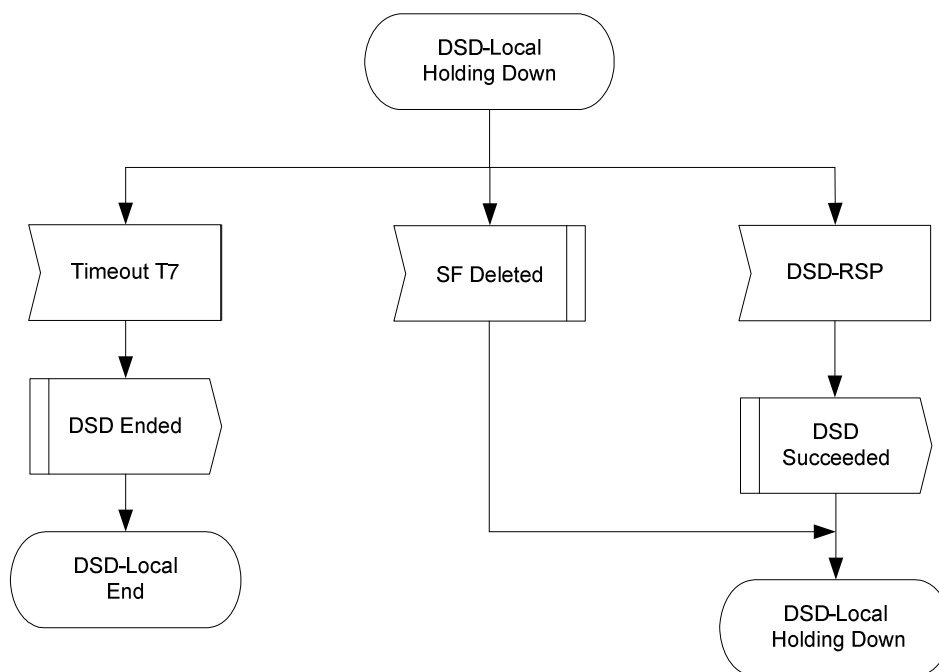


Figure 11-36: DSD-Locally Initiated Transaction Holding Down State Flow Diagram




Figure 11-37: DSD-Remotely Initiated Transaction Begin State Flow Diagram

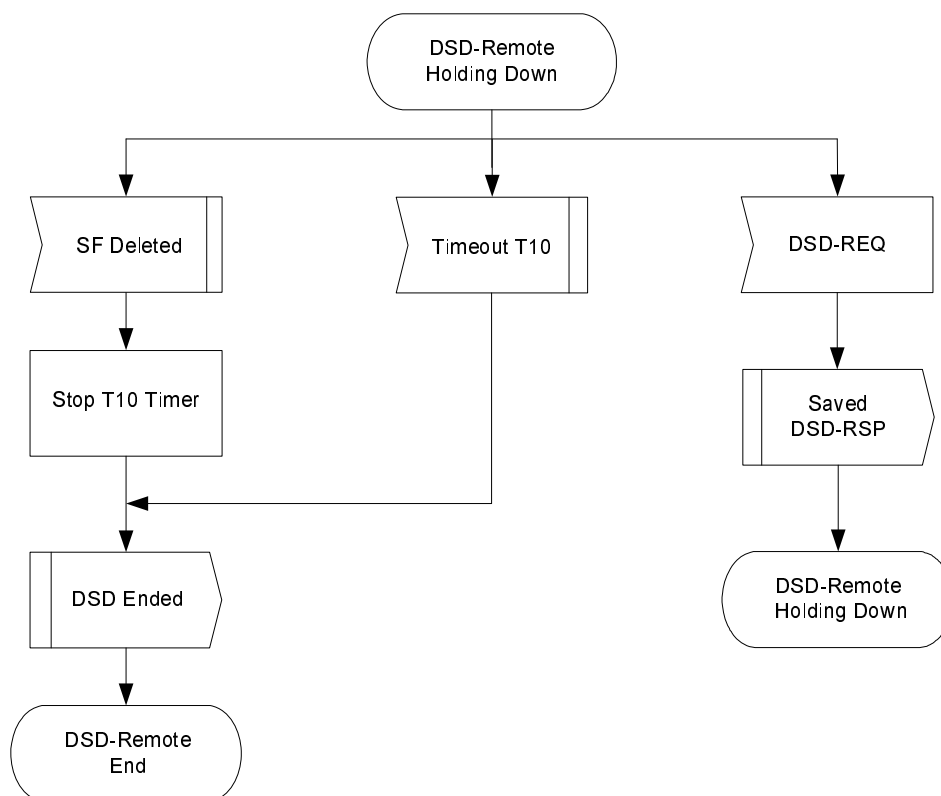


Figure 11-38: DSD-Remotely Initiated Transaction Holding Down State Flow Diagram

11.3 Pre-3.0 DOCSIS Upstream Channel Changes

At any time after registration, the CMTS may direct a Pre-3.0 DOCSIS CM to change its upstream channel with a UCC-REQ. This may be done for traffic balancing, noise avoidance or any of a number of other reasons which are beyond the scope of the present document. Figure 11-39 shows the procedure that MUST be followed by the CMTS. Figures 11-40 and 11-41 show the corresponding procedure at the CM.

The CMTS MUST NOT send a UCC-REQ to a docsis 3.0 CM. The equivalent function in DOCSIS 3.0 is achieved with a DCC-REQ or a DBC-REQ.

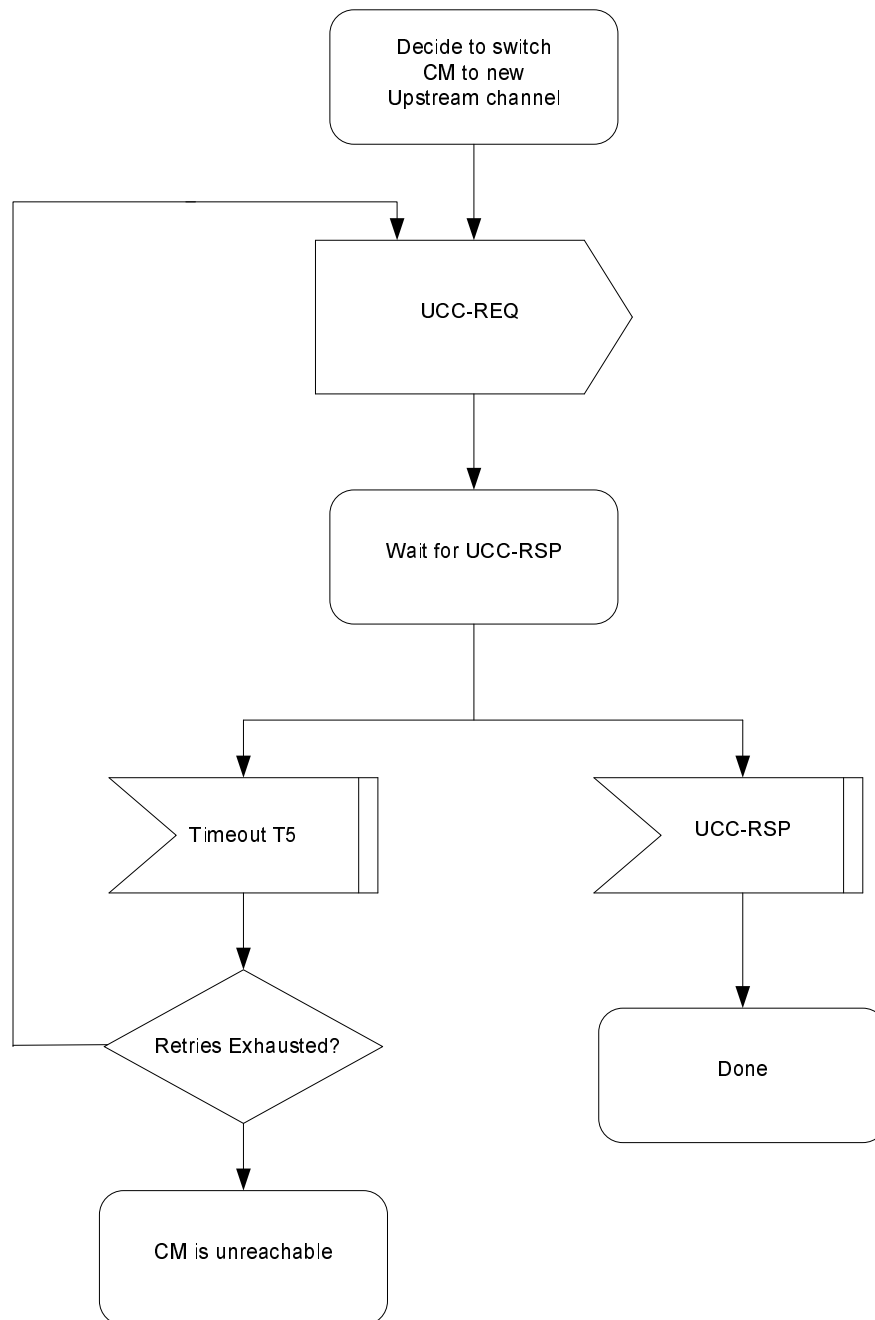
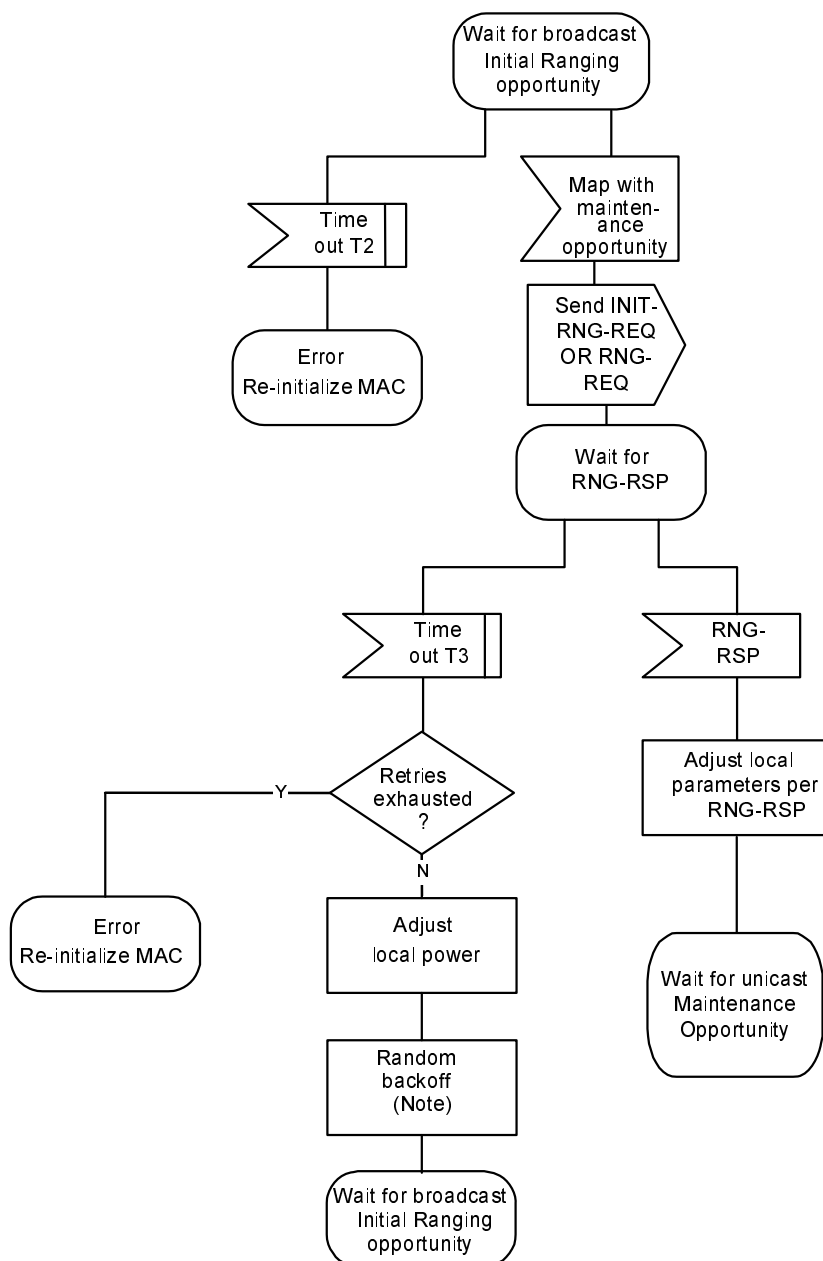


Figure 11-39: Changing Upstream Channels: CMTS View

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.



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

Figure 11-40: Changing Upstream Channels: CM View Part 1

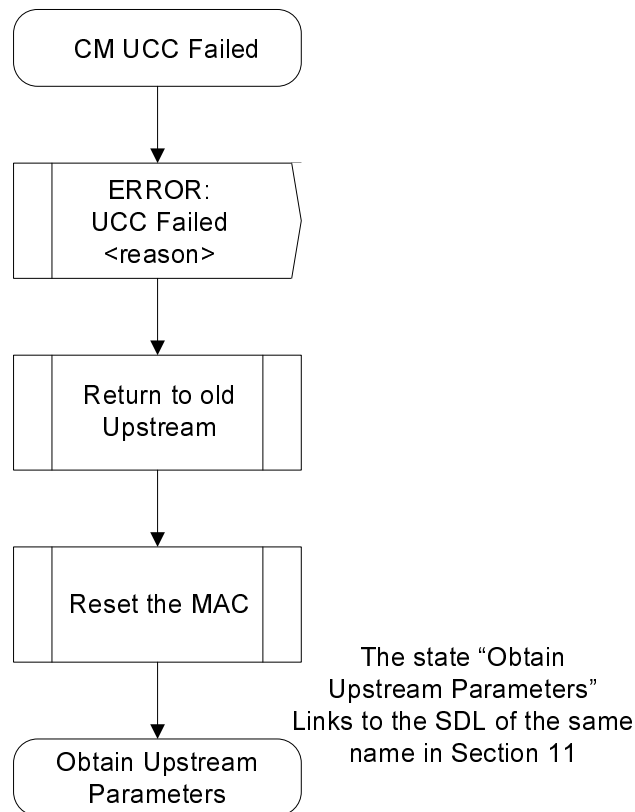


Figure 11-41: Changing Upstream Channels: CM View Part 2

When performing a requested Upstream Channel Change, upon synchronizing with the new upstream channel, the CM **MUST** perform broadcast initial ranging 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.

When performing a requested Upstream Channel Change, the CM **MUST NOT** perform re-registration, since its provisioning and MAC domain remain valid on the new channel.

If a CM in which DOCSIS 2.0 is enabled in registration is moved from a Type 1 channel to a Type 2 channel via a UCC, the CM **MUST** operate in 2.0 mode on the destination channel, basing its requests on IUCs 9 and 10 instead of IUCs 5 and 6. If a CM in which DOCSIS 2.0 is disabled in registration is moved from a Type 1 channel to a Type 2 channel via a UCC, the CM **MUST** continue to base its requests on the destination channel on IUCs 5 and 6.

11.4 Dynamic Downstream and/or Upstream Channel Changes

11.4.1 DCC General Operation

The Dynamic Channel Change (DCC) mechanism is intended for the following situations:

- changing the downstream channel and/or upstream channel of a CM not operating in Multiple Receive Channel mode;
- changing the MAC Domain of a CM operating in both Multiple Receive Channel mode and Multiple Transmit Channel Mode; and
- changing the upstream channel of a CM which was not assigned a Transmit Channel Configuration in the registration process and is thus not operating in Multiple Transmit Channel mode.

At any time after registration, the CMTS MAY use the DCC-REQ message to direct a CM not operating in Multiple Receive Channel mode to change its upstream and/or downstream channel. At any time after registration, the CMTS may use the DCC-REQ message to direct a CM to which a Transmit Channel Configuration was not assigned in the registration process to change its upstream channel. The CMTS MUST be capable of performing DCC operations to trigger upstream and/or downstream channel changes within a MAC domain and between MAC domains for CMs not operating in Multiple Receive Channel mode. The CMTS MUST be capable of performing DCC operations to a CM operating in Multiple Receive Channel mode to force it to reinitialize in a different MAC Domain. The CMTS MUST be capable of performing DCC operations to a CM which has not received a Transmit Channel Configuration in the registration process to force it to change its upstream channel. For a CM operating in Multiple Receive Channel mode, the CMTS will use Dynamic Bonding Change (DBC) messaging to change downstream channels within a MAC domain. For a CM to which a Transmit Channel Configuration was assigned in the registration process, the CMTS uses Dynamic Bonding Change (DBC) messaging to change upstream channels within a MAC domain.

Physical layer conditions permitting, the CMTS MUST be capable of executing Dynamic Channel Changes using all Initialization Techniques for CMs not operating in Multiple Receive Channel mode (see clause 11.4.1.2). This may be done for load balancing (as described in clause 11.6), noise avoidance or other reasons that are beyond the scope of the present document. In addition, the CMTS MUST support DCC operations triggered via SNMP (see [9]). Figure 11-42 through figure 11-45 show the procedure that MUST be followed by the CMTS when performing a dynamic channel change. Figure 11-46 through figure 11-49 show the corresponding procedure that MUST be followed by a CM when performing a dynamic channel change.

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 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.

When moving a CM within a MAC domain or when moving a CM to a new MAC domain with initialization technique other than zero, the Upstream Channel ID MUST be different 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 reinitialization, only ranging or neither. If this TLV is not present in DCC-REQ, the CM MUST reinitialize its MAC on the new channel assignment. (Refer to clause 11.4.1.2.) If the CM has been instructed to reinitialize, 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, a reinitialization may not be required. If the CM is being moved between MAC domains, a reinitialization may be required. Reinitializing, if requested, is done with the new upstream and downstream channel assignments. It includes obtaining upstream parameters, establish IP connectivity, establish time of day, transfer operational parameters, register and initialize baseline privacy. If reinitialization is performed, the CM MUST NOT send a DCC-RSP on the new channel.

As required in clause 6.4.20.1.3, if the CMTS sends a DCC-REQ to change the downstream of a CM operating in Multiple Receive Channel Mode, the CMTS will specify the Initialization Technique TLV that will reinitialize the CM MAC. As required in clause 6.4.20.1.3, if the CMTS sends a DCC-REQ to change the upstream of a CM to which a Transmit Channel Configuration was assigned in the registration process, the CMTS will specify the Initialization Technique TLV that will reinitialize the CM MAC. If the CMTS sends a DCC-REQ to change the upstream of a CM to which a Transmit Channel Configuration was not assigned in the registration process, the CMTS is not required to specify Initialization Technique 0 (reinitialize the MAC).

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 modulation rate, modulation type or mini-slot size have changed.

When DCC-REQ does not involve reinitialization 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 old 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 (depart) 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 reinitialize the MAC, then QoS resources are not reserved on the destination CMTS and T15 does not apply. If in the process of a dynamic channel change with a non-zero initialization technique the CMTS detects that the CM has reinitialized the MAC before completing the channel change, the CMTS MAY de-allocate the resources that were previously allocated to the modem on the new channel before the expiration of T15.

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) included in the DCC-REQ message and the local configuration of the destination CMTS.

The destination CMTS SHOULD calculate and limit T15 based on internal policy according to the guidelines in clause 11.4.1.1.

If initialization technique 1 (broadcast initial ranging) is utilized and if the CM arrives after T15 has passed and attempts 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 CM to reinitialize MAC.

When a CM is moved between downstream channels on different IP subnets using initialization techniques other than technique 0 (reinitialize 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 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.

11.4.1.1 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.

- $T15 = C_{mJumpTime} + C_{mtsRxRngReq}$.

Each of the variables in the equation calculating the T15 timer is explained in table 11-1.

Table 11-1: Variables Used to Calculate the T15 Timer

Variable	Explanation and Value
CmJumpTime	<p>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 the CM and returned in DCC-RSP (depart).</p> <p>If the CM does not specify the Jump Time TLV's, then the destination CMTS assumes that the value is 1,3 s. This recognizes the fact that the CM may continue to use the old channel until the expiry of the T13 timer.</p> <p>If the CM specifies the Jump Time TLV's, then the destination CMTS uses the specified value.</p>
CmtsRxRngReq	<p>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.</p> <p>For initialization technique 4 (Use Directly), this value is two times the CMTS time period between unicast station maintenance opportunities plus 20 milliseconds to 40 milliseconds for MAP and RNG-REQ transmission time and CMTS RNG-REQ processing time.</p> <p>For initialization technique 2 (unicast ranging), this value is two times the CMTS time period between unicast ranging opportunities plus 20 milliseconds to 40 milliseconds for MAP and RNG-REQ transmission time and CMTS RNG-REQ processing time.</p> <p>For initialization technique 1 (broadcast initial ranging), this value is 30 s. Because the variables involved in initial maintenance are not strictly under the control of the CMTS, the computation of this factor is uncertain.</p>

The minimum value of the T15 timer is four seconds; this was derived by quadrupling the value of the T13 timer. The maximum value of the T15 timer is 35 s.

11.4.1.2 Initialization Technique

There are many factors that drive the selection of an initialization technique when commanding a dynamic channel change. While it is desirable to provide the minimum of interruption to existing QoS services such as voice over IP or streaming video sessions, a CM will not be able to successfully execute a channel change given an initialization technique that is unsuitable for the cable plant conditions. A CM may impair the new channel if it is commanded to use an unsuitable initialization technique. For instance, consider the use of initialization technique 4 (Use Directly) for a DCC changing an upstream channel when there is a significant difference in propagation delay between the old and new upstream channel. Not only will the CM be unable to communicate with the CMTS after the channel change, but its transmissions may also interfere with the transmissions of other CMs using the channel.

Careful consideration needs to be given to the selection of an initialization technique. Some restrictions are listed below. This list is not exhaustive, but is intended to prevent the use of a particular initialization technique under conditions where its use could prevent the CM from successfully executing the channel change. Packets may be dropped under some conditions during channel changes; 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.

If a CM performs a channel change without performing a MAC reinitialization (as defined in clause 6.4.20.1.3), all the state variables of the CM MUST remain constant (BPI keys, IP address, Classifiers, PHS Rules, etc.), with the exception of the state variables which are explicitly changed by the DCC-REQ message encodings (as defined in clause 6.4.20.1.1 through clause 6.4.20.1.7). 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. Note that regardless of the initialization technique, the CPE will not be aware of any configuration changes and will continue to use its existing IP address.

11.4.1.2.1 Initialization Technique Zero (0)

The use of initialization technique 0 (reinitialize the MAC), results in the longest interruption of service. The CMTS MUST signal the use of this technique when QoS resources will not be reserved on the new channel(s), when the downstream channel of a CM confirmed with Multiple Receive Channel Support is changed or when the upstream channel of a CM to which a Transmit Channel Configuration was assigned in the registration process is changed. The CMTS MUST use initialization technique 0 in DCC messages to CMs operating in Multiple Transmit Channel mode and Multiple Receive Channel mode.

11.4.1.2.2 Initialization Technique One (1)

The use of initialization technique 1 (broadcast initial ranging) may also result in a lengthy interruption of service. However, this interruption of service is mitigated by the reservation of QoS resources on the new channel(s). The service interruption can be further reduced if the CMTS supplies downstream parameter sub-TLV's and the UCD substitution TLV in the DCC-REQ in addition to providing more frequent initial ranging opportunities on the new channel.

11.4.1.2.3 Initialization Technique Two (2)

The use of initialization technique 2 (unicast ranging) offers the possibility of only a slight interruption of service. In order to use initialization technique 2, the CMTS MUST:

- Synchronize timestamps (and downstream symbol clocks for S-CDMA support) across the downstream channels involved and specify SYNC substitution sub-TLV with a value of 1 if the downstream channel is changing.
- Include the UCD substitution in the DCC message if the upstream channel is changing.

However, the CMTS MUST NOT use initialization technique 2 if:

- The DCC-REQ message requires the CM to switch between S-CDMA and TDMA.
- Propagation delay differences between the old and new channels will cause the CM burst timing to exceed the ranging accuracy requirements of [12].
- Attenuation or frequency response differences between the old and new upstream channels will cause the received power at the CMTS to be outside the limits of reliable reception.

11.4.1.2.4 Initialization Technique Three (3)

The use of initialization technique 3 (initial ranging or periodic ranging) offers the possibility of only a slight interruption of service. 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. However, the CMTS MUST NOT use initialization technique 3 if the conditions for using techniques 1 and 2 are not completely satisfied.

11.4.1.2.5 Initialization Technique Four (4)

The use of initialization technique 4 (use the new channel without reinitialization or ranging) results in the least interruption of service.

In order to use initialization technique 4, the CMTS MUST:

- Synchronize timestamps (and downstream symbol clocks for S-CDMA support) across the downstream channels involved and specify SYNC substitution sub-TLV with a value of 1 if the downstream channel is changing.
- Include the UCD substitution in the DCC message if the upstream channel is changing.

However, the CMTS MUST NOT use initialization technique 4 if:

- The CM is operating in S-CDMA mode and any of the following parameters are changing:
 - Modulation Rate.
 - S-CDMA US ratio numerator 'M'.
 - S-CDMA US ratio denominator 'N'.
 - Downstream channel.
- The DCC-REQ message requires the CM to switch between S-CDMA and TDMA.

- Propagation delay differences between the old and new channels will cause the CM burst timing to exceed the ranging accuracy requirements of [12].
- Attenuation or frequency response differences between the old and new upstream channels will cause the received power at the CMTS to be outside the limits of reliable reception.
- Micro-reflections on the new upstream channel will result in an unacceptable PER (greater than 1 %) with the pre-equalizer coefficients initialized according to [12].

11.4.2 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 annex C) 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 can satisfy a CMTS-originated service flow add or change (e.g. for PacketCable Multimedia) on a different downstream or upstream channel for a CM not operating in Multiple Transmit Channel mode or Multiple Receive Channel mode, the CMTS SHOULD execute the DCC command first and then issue a DSA or DSC command to that CM.

If the provisioning system default is to specify the upstream channel ID, the downstream frequency and/or a downstream channel list in the configuration file, care should be taken when using DCC, particularly when using initialization technique 0 (reinitialize MAC). If a CMTS does a DCC with reinitialize, the config file could cause the CM to come back to the original channel. This would cause an infinite loop.

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 DCC command if the CMTS has previously issued a DBC command and that command is still outstanding.

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 annex C). 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 reinitialize its MAC. The previous channel assignment represents a known good operating point which should speed up the reinitialization 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" (annex 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 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" (annex 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 annex C).

If a DOCSIS 3.0 CM that has DOCSIS 2.0 Mode enabled is moved from a Type 1 channel to a Type 2 or Type 4 channel via a DCC using an initialization technique other than 0 (reinitialize MAC), the CM MUST base its requests on IUCs 9 and 10 instead of IUCs 5 and 6 (i.e. the CM operates in DOCSIS 2.0 Mode on the destination channel). If a CM in which DOCSIS 2.0 is disabled in registration is moved from a Type 1 channel to a Type 2 channel via a DCC with initialization technique other than 0 (reinitialize MAC), the CM MUST continue to base its requests on the destination channel on IUCs 5 and 6. A CM in which DOCSIS 2.0 is disabled in registration MUST consider a Type 3 or a Type 4 channel to be invalid for any DCC using initialization technique other than 0 (reinitialize MAC).

If the CM is moved via a DCC using the initialization technique of 0 (reinitialize MAC), the previous Enable 2.0 Mode setting is not applicable. If DOCSIS 2.0 Mode was previously disabled and the target upstream channel is Type 2, Type 3 or Type 4, the 2.0 Mode disable setting is discarded and the CM MUST use the target upstream channel. However, after reinitialization, the CM MUST operate using the Enable 2.0 Mode setting received in the configuration file acquired in the reinitialization process.

11.4.3 DCC State Transition Diagrams

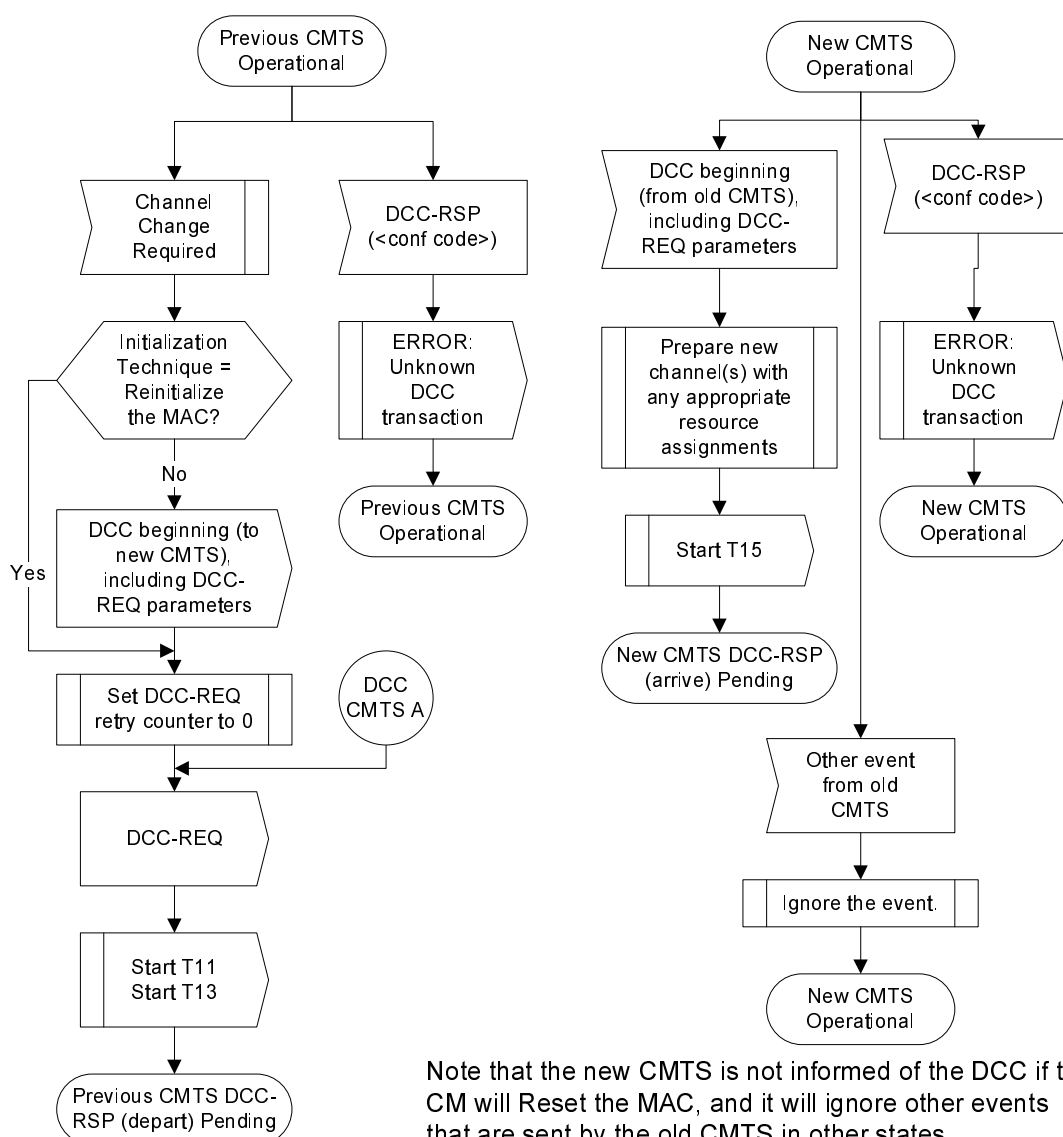


Figure 11-42: Dynamically Changing Channels: CMTS View Part 1

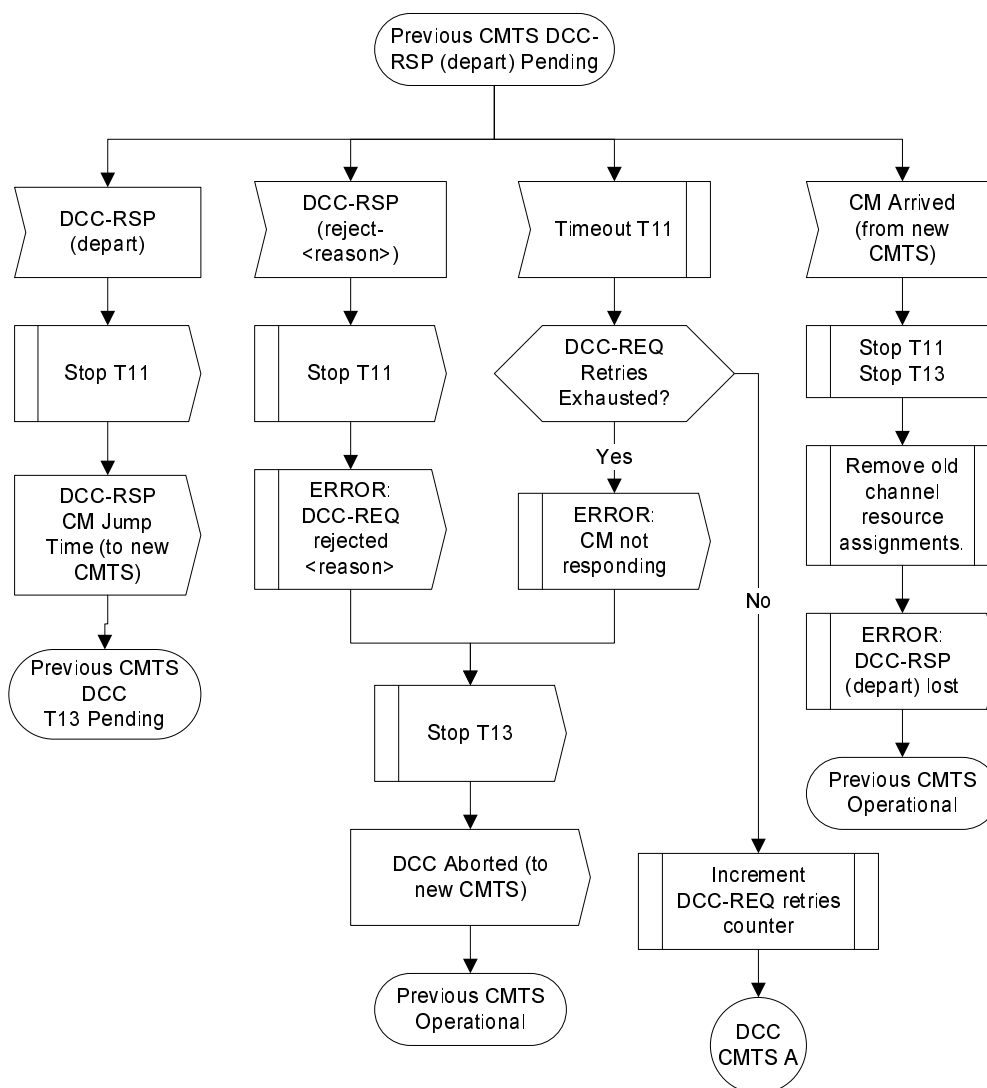


Figure 11-43: Dynamically Changing Channels: CMTS View Part 2

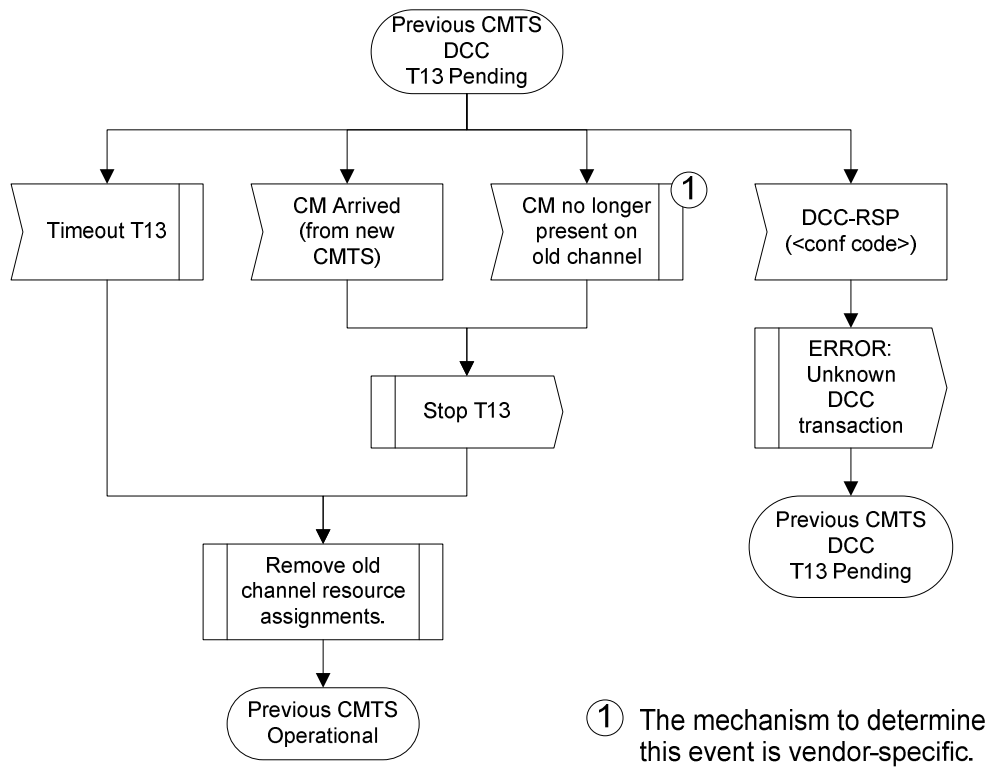


Figure 11-44: Dynamically Changing Channels: CMTS View Part 3

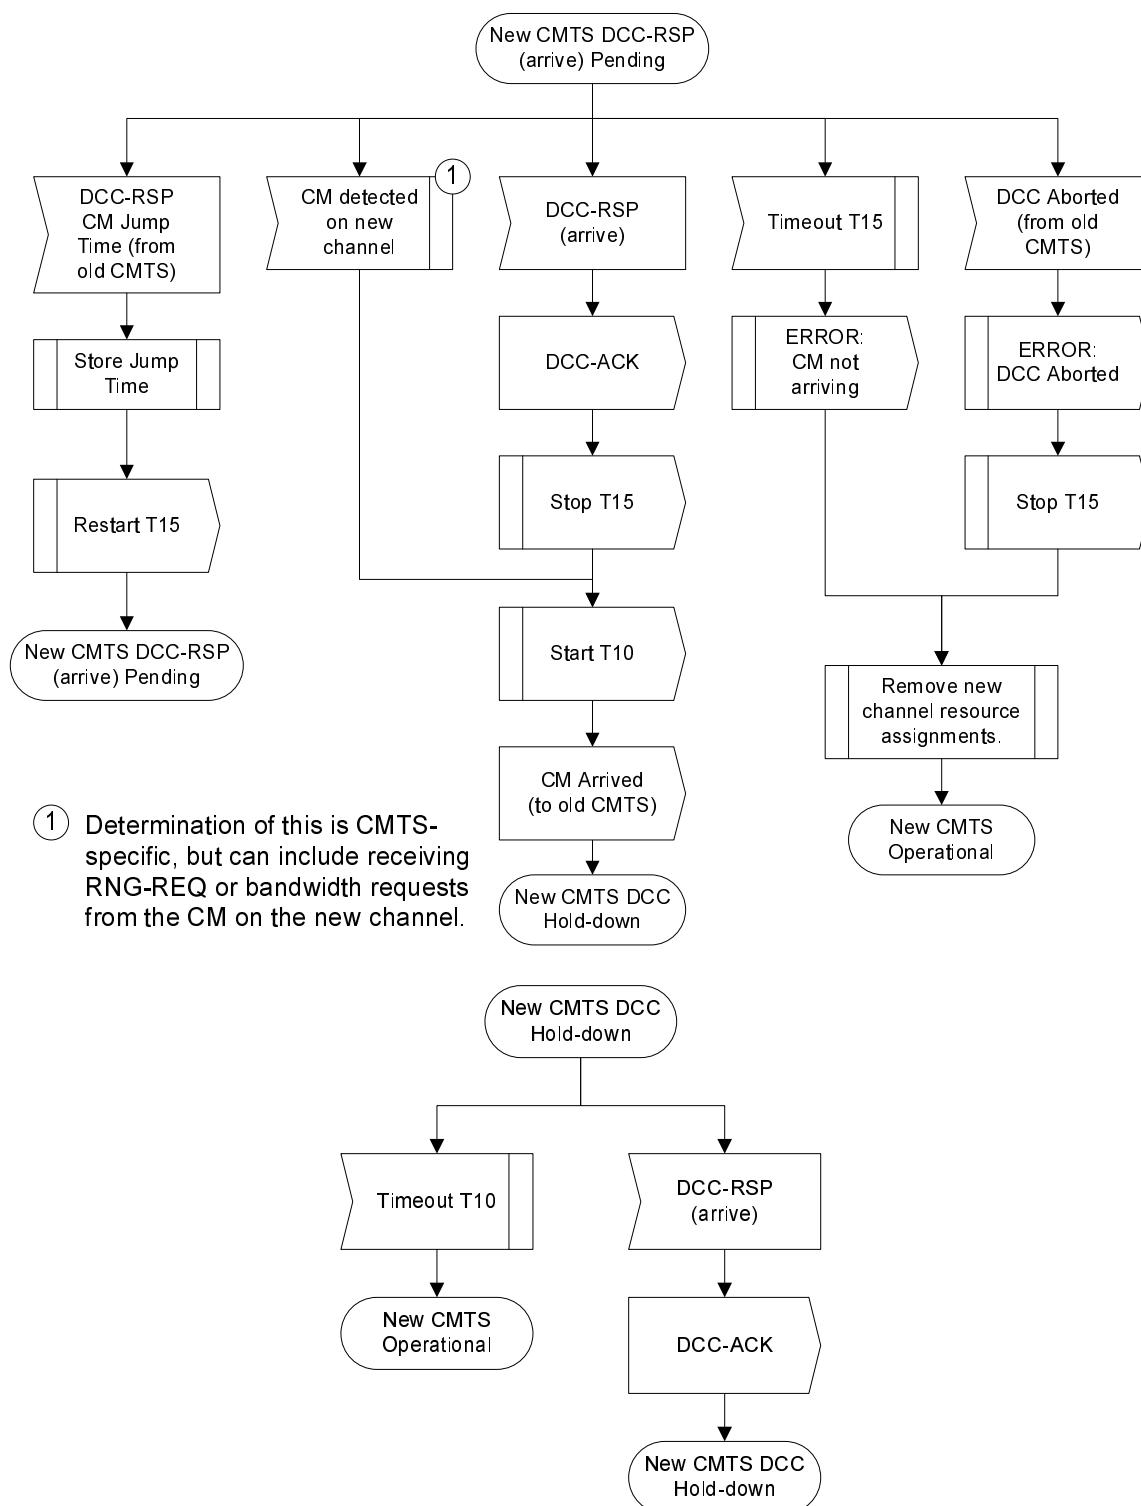


Figure 11-45: Dynamically Changing Channels: CMTS View Part 4

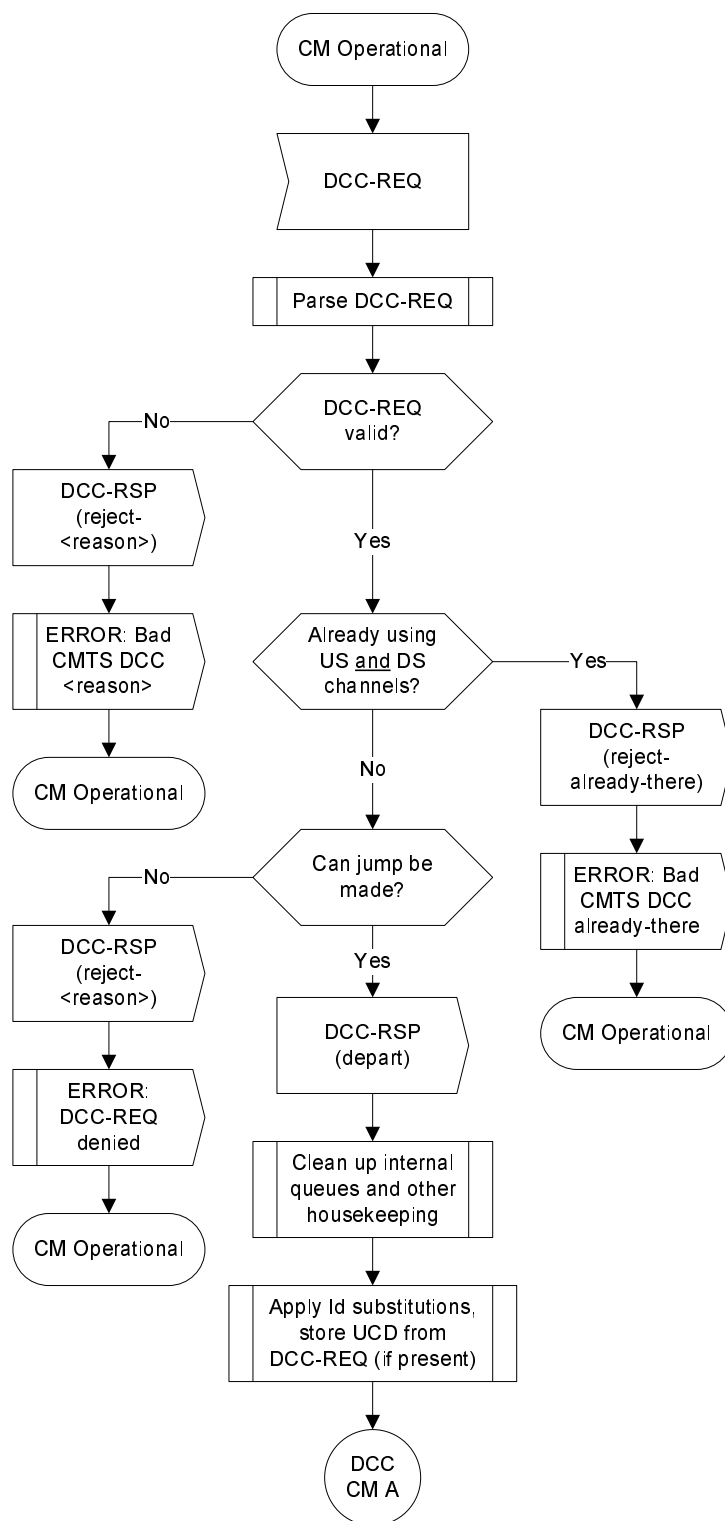


Figure 11-46: Dynamically Changing Channels: CM View Part 1

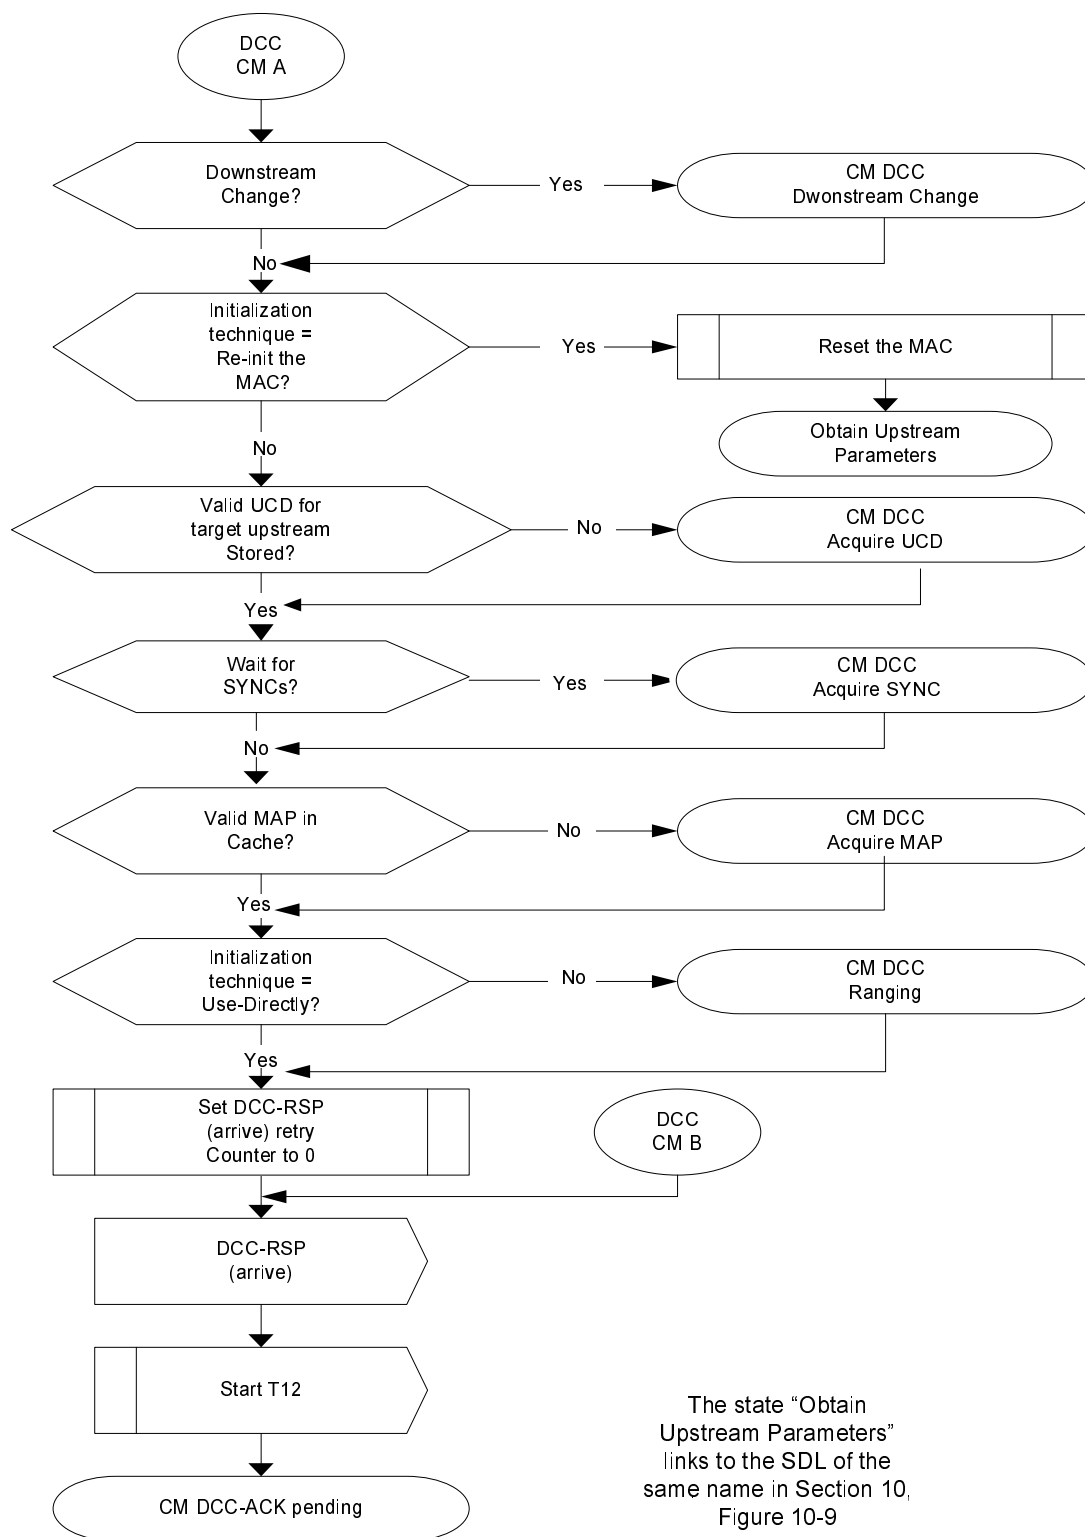


Figure 11-47: Dynamically Changing Channels: CM View Part 2

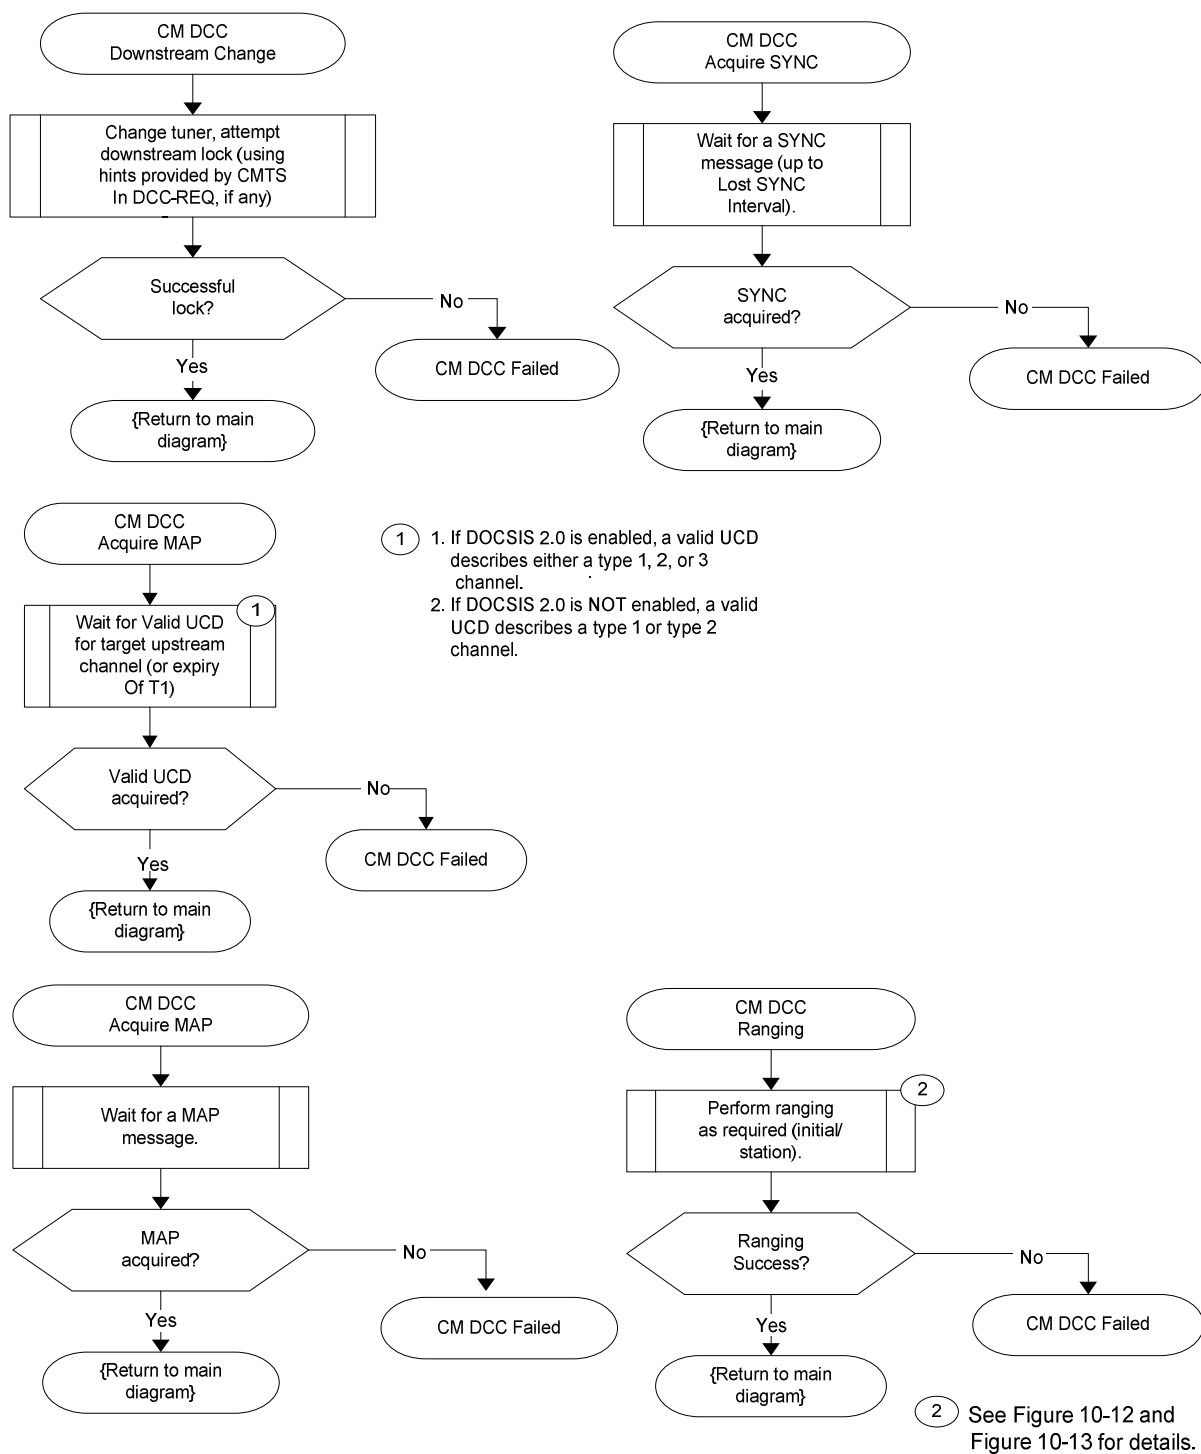


Figure 11-48: Dynamically Changing Channels: CM View Part 3

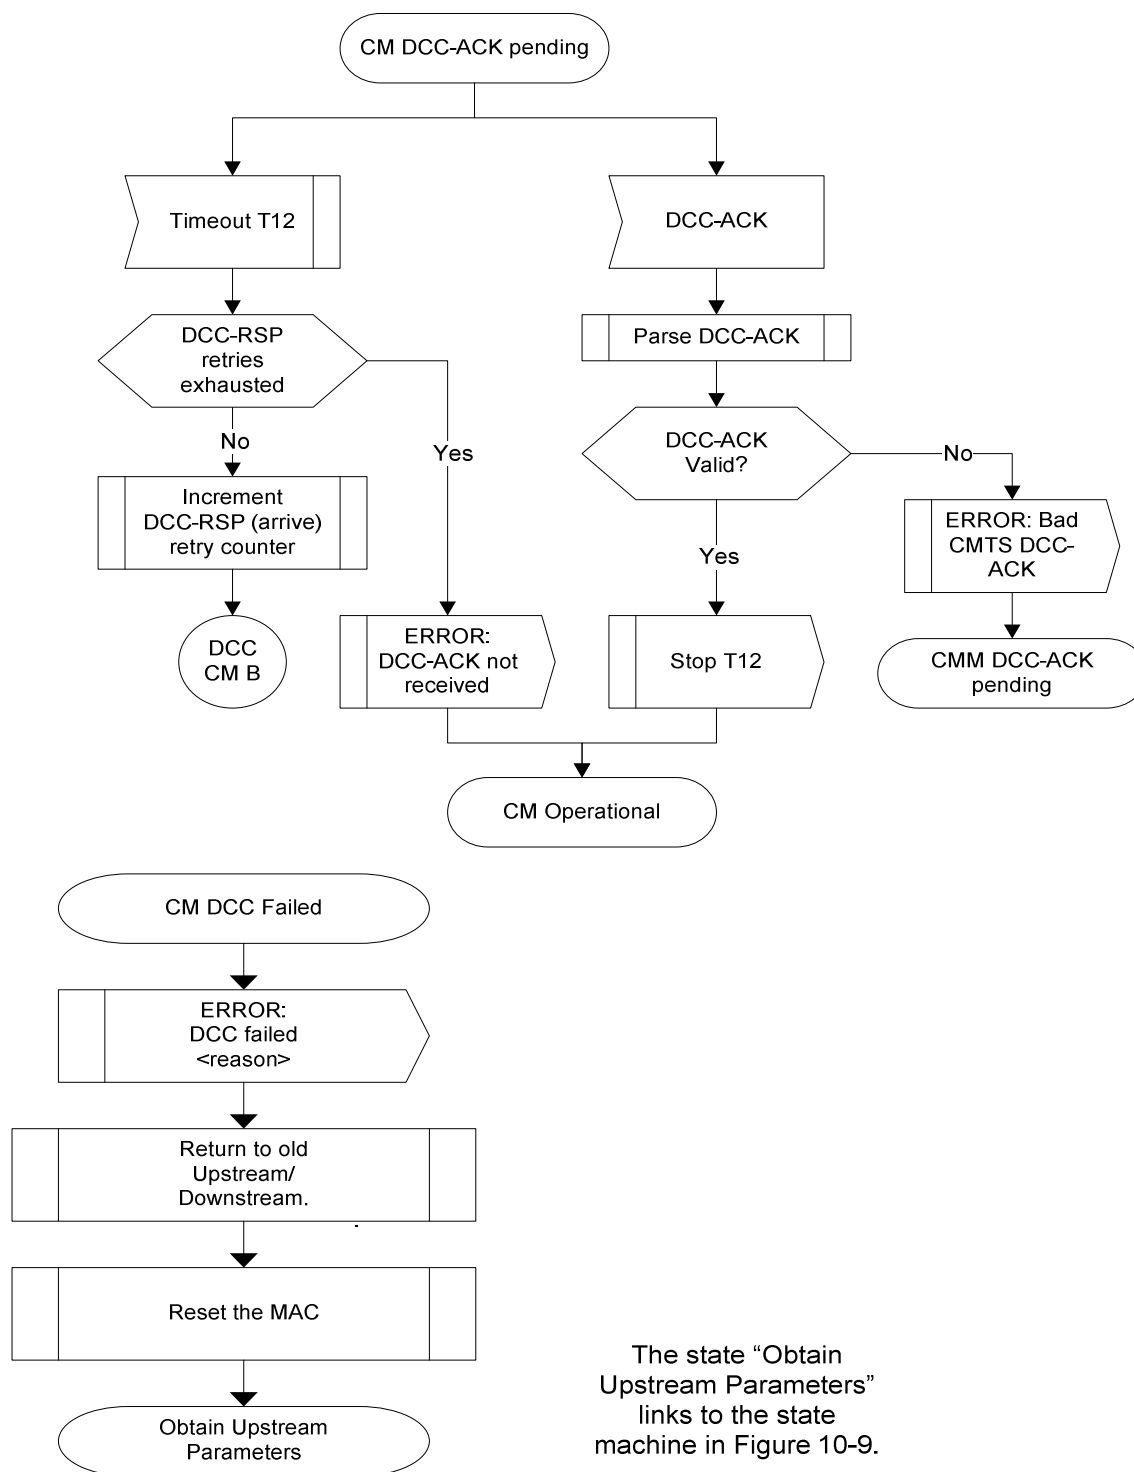


Figure 11-49: Dynamically Changing Channels: CM View Part 4

11.4.4 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 housecleaning, 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 table 11-2.

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.

Table 11-2: Factors affecting DCC performance

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 Broadcast Initial Ranging	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 Unicast Ranging	The DCC transaction SHOULD complete within 1,5 s after the start of jump.
	3 Broadcast or Unicast Ranging	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.
DS Parameter		The CMTS SHOULD include the downstream parameter TLV's for station maintenance and use directly initialization techniques that are expected to occur quickly.
UCD Substitution		The CMTS MUST set the UCD substitution TLV according to clause 11.4.1.2.
SYNC Substitution		The CMTS MUST set the SYNC substitution TLV according to clause 11.4.1.2.
CM Jump Time		The length of jump TLV SHOULD be less than one second for downstream channel changes that include the downstream parameter TLV's or for upstream only channel changes.

When the DCC-REQ does not contain UCD Substitution TLV's 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 TLV's 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.

11.5 Dynamic Bonding Change (DBC)

11.5.1 DBC General Operation

At any time after registration, the CMTS uses the DBC command to change any combination of the following parameters in a CM:

- The receive channel set.
- DSID(s) or DSID associated attributes.
- Security association(s) for encrypting downstream traffic.
- The transmit channel set.
- Service Flow Cluster Assignments.

The CMTS MUST be capable of performing DBC operations within a MAC domain. The DBC change can only occur within a MAC domain; the CMTS moves the CM between MAC domains using the DCC message. The CMTS MUST NOT initiate a DBC transaction to direct any of the CM's channels to a different MAC domain.

Multiple actions can occur within a single DBC message. If the DBC-REQ contains a change in the RCS, the CM MUST implement the downstream channel changes prior to making any other changes in the DBC-REQ message.

11.5.1.1 Changes to the Receive Channel Set

The CMTS can add channels to the Receive Channel Set, delete channels from the Receive Channel Set or change channels within the Receive Channel Set of a CM by sending the CM a new Receive Channel Configuration via a DBC-REQ (see clause C.1.5.3). If the CM receives a DBC-REQ with a Receive Channel Configuration that the CM is not capable of using, the CM MUST reject the DBC-REQ. A Receive Channel Set is the complete list of all the Downstream Channels that were included in the RCC of the DBC exchange.

Changes in the RCC that affect the CM's Primary Downstream Channel will require the CM to re-range on its upstream channels before it can continue operation. Specifically, changes to the Primary Downstream Channel itself or changes to the Receive Module(s) to which the Primary Downstream Channel is connected (either directly or indirectly) will require the CM to re-range. If the CMTS makes a change in the CM's RCC that affects the CM's Primary Downstream Channel, the CMTS MUST signal re-ranging and include an initialization technique in the DBC-REQ for all upstream channels. This means that the CMTS cannot make changes affecting the Primary Downstream Channel using DBC unless a TCC encoding has been included in the REG-RSP-MP. If the CMTS does not include an initialization technique for each upstream channel in the transmit channel set in the DBC-REQ when the CM's primary downstream is affected, the CM MUST reject the DBC-REQ message.

Clause 11.5.3 details the operation of the CMTS and CM during the DBC process. With the exception of a complete change to the receive channel set, the CMTS stops sending traffic on any channels to be deleted from the RCS. When removing channels from the RCS, the CMTS has several means of minimizing packet loss. The CMTS may choose to stop sending traffic on the downstream channels to be removed from the RCS prior to sending the DBC-REQ message. The CMTS may use a vendor-specific delay between control and data messages. In addition, the CMTS may transmit the DBC-REQ messages on the highest latency downstream channel. In the case of a complete change to the RCS, stopping traffic on the original RCS could cause an interruption in traffic that could persist for some time in the event of a lost DBC-REQ message. In this case, the CMTS has the option of duplicating traffic on the old and new RCS. The CMTS sends the DBC-REQ and waits for the DBC-RSP. Once the CMTS receives the DBC-RSP, it begins transmitting packets on the new channel set. The CMTS waits a vendor-specific time before sending the DBC-ACK to account for differences in delay between control messages and data messages and to ensure that the CM receives all data traffic sent prior to the DBC-ACK message. The CMTS then sends a DBC-ACK.

When the CM receives an invalid DBC-REQ, the CM sends a DBC-RSP rejecting the message. The CM MUST include an applicable error encoding in the DBC-RSP for at least one top-level TLV in the DBC-REQ that the CM could not implement. If the DBC-REQ message is valid, the CM then makes the receive channel set changes identified in the DBC-REQ. Once the CM has completed all attempts to acquire all new channels and deletes any channels being removed from the RCS, the CM sends a DBC-RSP and waits for a DBC-ACK. The CM MUST try acquisition of the new DS channels in the RCS for the duration of the T(dbc downstream acquisition timer) before declaring that it was unable to acquire the downstream channel; downstream acquisition consists of QAM lock, FEC lock and synchronization of MPEG framing. The DBC-RSP will contain no errors if the CM was able to make all RCS changes, Partial Service errors if the CM was able to make some of the RCS changes or failure if the CM was unable to make any RCS changes. When the CM receives a DBC-ACK, the CM enables rapid loss detection of all resequencing DSIDs.

If the CMTS does not receive the DBC-RSP after all the retries and the RCS changed, the CMTS either returns the traffic to the previous downstream channel(s) or stops duplicating traffic after the expiration of the Initializing Channel Timeout timer, depending on previous operation. If the CMTS does not receive the DBC-RSP after all the DBC-REQ retries and the RCC contained a change that affected the CM's primary downstream, the CMTS reinitializes the CM using the CM-CTRL-REQ message. The CMTS MUST send the CM-CTRL-REQ message on either an overlapping downstream channel or if there were no overlapping channels, on both the old and new channels to ensure that the CM received the message. The CMTS also has the option of discontinuing station maintenance for all upstream channels associated with the CM to ensure that a reinitialization occurs. If the CMTS does not receive the DBC-RSP after all the DBC-REQ retries and the new RCC did not contain a change that affected the CM's primary downstream, the CMTS MUST recover from this condition. Recovery is considered complete if the CM's receive channel set is synchronized at both the CMTS and CM. The CMTS actions may include the following for a RCS replacement:

- initiation of a new DBC transaction to retry the errored DBC transaction;
- initiation of a new DBC transaction to undo the errored DBC transaction; or
- reinitialization of the CM using the CM-CTRL-REQ message.

In order for the CM to complete the DBC, it is necessary that the CM be able to tune to at least the Primary Downstream Channel. If the CM cannot tune to the new Primary Downstream Channel included the RCS, the CM MUST reinitialize the MAC and return to the previous primary downstream. If the CM tunes to the Primary Downstream Channel, but cannot tune to all of the new channels in the RCS, the CM logs an error and MUST send a DBC-RSP with partial service. In this case, the CM goes operational on the channels on which it is able to tune. The CMTS may attempt to remedy any partial service state by any combination of the following:

- initiating a new DBC transaction to add missing channels;
- reinitializing the CM; or
- moving the CM to a different set of channels.

11.5.1.2 Changes to a DSID

Using DBC messaging, the CMTS can change attributes of a DSID. DSID attributes that can change are:

- Resequencing Encodings:
 - Downstream Resequencing Channel List.
 - DSID Resequencing Wait Time.
- Resequencing Warning Threshold.
- CM-STATUS Hold-Off Timer for Out-of-range Events.
- Multicast Encodings:
 - Client MAC Address.
 - Multicast CM Interface Mask.
 - Group MAC Address.

11.5.1.2.1 Changes to Resequencing Encodings

11.5.1.2.1.1 Changes to the Downstream Resequencing Channel List

The CMTS can add channels to the Downstream Resequencing Channel List, delete channels from the Downstream Resequencing Channel List or change channels within the Downstream Resequencing Channel List by replacing the CM's Downstream Resequencing Channel List with a new Downstream Resequencing Channel List.

Clause 11.5.3 details the operation of the CMTS and CM during the DBC process. When no RCS changes are required, the CMTS implements changes to the Downstream Resequencing Channel List by continuing to transmit packets over the old Downstream Resequencing Channel List when sending the DBC-REQ message until the DBC-RSP message confirms that the CM has accepted the new Downstream Resequencing Channel List. Once the CMTS receives the DBC-RSP, it begins transmitting packets with the associated resequencing DSID on the new Downstream Resequencing Channel List. The CMTS waits a vendor-specific time before sending the DBC-ACK to account for differences in delay between control messages and data messages and to ensure that the CM receives all data traffic sent prior to the DBC-ACK message. The CMTS then sends a DBC-ACK.

When the CM receives the DBC-REQ, it expands the rapid loss detection of a resequencing DSID across the union of the old Downstream Resequencing Channel List and the new Downstream Resequencing Channel List and sends a DBC-RSP. The CM also expands its filters such that it discards packets received on a downstream channel not included in this union. When the CM receives a DBC-ACK, the CM waits the duration of the DSID Resequencing Wait Time and contracts the rapid loss detection to the new Downstream Resequencing Channel List. The CM then discards any DSID-labeled frames received on downstream channels not in the new Downstream Resequencing Channel List.

If the Downstream Resequencing Channel List changed with no changes to the RCS and the CMTS does not receive the DBC-RSP after all the retries, the CMTS MUST return traffic associated with the Resequencing DSID to the previous Downstream Resequencing Channel List. If the Downstream Resequencing Channel List changed with no changes to the RCS and CMTS does not receive the DBC-RSP after all the retries, the CMTS MUST recover from this condition. Recovery is considered complete if the CM's Downstream Resequencing Channel List is synchronized at both the CMTS and CM. The CMTS actions may include the following for a Downstream Resequencing Channel List replacement without an RCS replacement:

- Initiation of a new DBC transaction to delete the DSID associated with the Downstream Resequencing Channel List.
- Initiation of a new DBC transaction to retry the errored DBC transaction.
- Initiation of a new DBC transaction to undo the errored DBC transaction.
- Reinitialization of the CM using the CM-CTRL-REQ message.

If the CM does not receive a DBC-ACK after all the retries, the CM logs an error, goes operational and restores rapid loss detection of the Resequencing DSID.

11.5.1.2.1.2 Changes to the DSID Resequencing Wait Time

Clause 11.5.3 details the operation of the CMTS and CM during the DBC process. Skew may change due to network changes in the CIN or other circumstances on the network (clause 8.2.3.1) even when no RCS or Downstream Resequencing Channel List changes occur. Further, configuration changes at the CMTS may also have an impact on skew. The CMTS may choose to communicate this change in skew to the CM via a change in the DSID Resequencing Wait Time.

If the CMTS has been requested to perform a reconfiguration that results in a reduction in skew, the CMTS SHOULD perform the reconfiguration prior to sending any DBC-REQ message communicating a change in the DSID Resequencing Wait Time to an affected modem. The CMTS SHOULD wait a vendor-specific time before sending the DBC-REQ to the first modem to account for differences in delay between control messages and data messages and to ensure that the CM receives all data traffic sent prior to the DBC-REQ message. After sending the DBC-REQ message, the CMTS waits for the DBC-RSP message. Once the CMTS receives the DBC-RSP message, the CMTS sends a DBC-ACK message.

If the CMTS has been requested to perform a reconfiguration that results in an increase in skew, the CMTS may choose to modify the DSID Resequencing Wait Time. If it modifies this parameter, the CMTS sends DBC-REQ messages to all affected modems and waits for DBC-RSP messages. The CMTS SHOULD perform the reconfiguration after the CMTS receives the DBC-RSP from all affected modems.

When the CM receives the DBC-REQ, it applies the change in the DSID Resequencing Wait Time. After it completes implementation of the modified DSID, the CM sends a DBC-RSP.

11.5.1.2.2 Changes to Multicast Encodings

The CMTS can initiate a DBC transaction to either add a multicast DSID, change attributes of an existing multicast DSID or delete a multicast DSID. Clause 11.5.3 details the operation of the CMTS and CM during the DBC process. When no RCS or Downstream Resequencing Channel List changes are required, the CMTS implements changes of some multicast DSID attributes prior to sending the DBC-REQ message and some changes after receipt of the DBC-RSP message. The CMTS sends the DBC-REQ message containing multicast encodings and waits for the DBC-RSP message. Once the CMTS receives the DBC-RSP, it sends a DBC-ACK.

Although the CMTS may forward multicast traffic labeled with the new or modified DSID at any time, the CM will not forward the packets labeled with the new or modified DSID until after it receives the DBC-REQ message containing the DSID. When the CMTS is required to start a new multicast replication for a CM joining a multicast session, the CMTS has the option of waiting to forward the multicast traffic to that CM until the DBC-RSP from the CM is received to ensure that the CM will not discard the multicast traffic because the DSID is not configured. Alternatively, the CMTS may forward multicast traffic with this DSID after sending the DBC-REQ but before receiving the DBC-RSP from the CM. By not waiting for the DBC-RSP from the CM, the CMTS can start the multicast traffic sooner which avoids any delays induced by waiting for the DBC-RSP.

When the CM receives the DBC-REQ, it implements the change in the multicast DSID attribute. After it completes implementation of the DSID modifications, the CM sends a DBC-RSP.

If the Multicast Encodings of a DSID are changed and CMTS does not receive the DBC-RSP after all the DBC-REQ retries, the CMTS does not know whether the CM has implemented the DBC or not. The CMTS MUST recover from this condition. Recovery is considered complete if the state of the Multicast DSID is synchronized at both the CMTS and the CM. The recommended CMTS recovery action for this condition is to initiate a new DBC transaction to delete the modified multicast DSID. Other CMTS actions may include the following:

- Initiation of a new DBC transaction to retry the errored DBC transaction.
- Initiation of a new DBC transaction to undo the errored DBC transaction.
- Reinitialization of the CM using the CM-CTRL-REQ message.

When the CM receives a DBC-REQ adding or changing a multicast DSID, the CM associates the client MAC address and Multicast CMIM encodings to a list of interfaces and forwards traffic to the appropriate interface accordingly.

When adding a multicast DSID, the CMTS MUST include a Client MAC Address Encoding and/or a Multicast CMIM in the DBC-REQ message:

- If the CMTS includes only Client MAC Address encodings, the CM MUST associate the interface(s) identified by the client MAC addresses with the DSID. The CM MUST assume that the multicast CMIM is all zeros for this DSID.
- If the CMTS includes only the Multicast CMIM, the CM MUST associate the interfaces provided in the Multicast CMIM with the DSID.
- If the CMTS includes both Client MAC Address Encodings and a Multicast CMIM, the CM MUST associate the union of the interfaces identified by the client MAC addresses and the interfaces identified in the Multicast CMIM with the DSID.

When changing attributes of an existing multicast DSID, any of the following combinations are valid:

- If the CMTS includes neither Client MAC Address Encodings nor Multicast CMIM for a particular DSID, the CM MUST keep the current association of interfaces with the DSID unchanged.
- If the CMTS includes only the Client MAC Address Encodings for a particular DSID, the CM updates the list of Client MAC Addresses according to the new Client MAC Address Encodings and keeps the CMIM unchanged. The CM MUST associate the union of the interface(s) identified by the updated client MAC address(es) and the interfaces identified in the current Multicast CMIM with the DSID.
- If the CMTS includes only the Multicast CMIM for a particular DSID, the CM updates the CMIM and keeps the list of Client MAC Addresses unchanged. The CM MUST associate the union of the interfaces of the current Client MAC Address(es) and the interfaces enabled by the new Multicast CMIM with the DSID.
- If the CMTS includes both the Client MAC Address Encodings and Multicast CMIM for a particular DSID, the CM updates the current list of Client MAC Addresses according to the new Client MAC Address Encodings and updates the CMIM. The CM MUST associate the DSID with the union of the interfaces identified by the updated Client MAC address list and the interfaces enabled by the new Multicast CMIM.

When deleting a multicast DSID, the Client MAC Address Encodings and Multicast CMIM are ignored by the CM. The CM deletes the DSID and all associated forwarding information.

The CMTS can remove Client MAC Addresses associated with a DSID in two ways. The CMTS can either send a DBC message to change the DSID with those Client MAC addresses deleted or the CMTS can send a DBC message that deletes the DSID.

11.5.1.3 Changes to the Security Association for encrypting downstream traffic

Using DBC messaging, the CMTS can add or delete Security Associations (SA) used to encrypt downstream traffic. The CMTS is not allowed send a DBC-REQ to a CM that is not in the "Authorized" State. The CMTS is allowed send a DBC-REQ with an SA that employs a cryptographic suite unsupported by the CM. If an unauthorized CM receives a DBC-REQ with a Security Association, the CM rejects the DBC-REQ. If the CM receives a DBC-REQ with a Security Association that the CM is not capable of using, the CM rejects the DBC-REQ [15].

Clause 11.5.3 details the operation of the CMTS and CM during the DBC process. When changes to the security associations for encrypting downstream traffic are necessary for multicast flows, the CMTS communicates the SA changes to the CM in a DBC-REQ message and waits for the DBC-RSP message. Once the CMTS receives the DBC-RSP, it sends a DBC-ACK.

When the CM receives a DBC-REQ adding an SA for which the CM is not already running a TEK state machine and the CM supports the cryptographic suite identified, the CM adds the SA and initiates a TEK state machine for the new SA. If the CM is already running a TEK state machine for the signaled SA or the CM does not support the cryptographic suite identified in the SA, the CM rejects the DBC-REQ. When the CM receives a DBC-REQ deleting an SA, it deletes the SA and terminates the associated TEK state machine for that SAID. The CM then sends a DBC-RSP and waits for a DBC-ACK.

Although the CMTS may start encrypting traffic with this SAID at any time, the CM will not forward the packets encrypted with this SAID until it completes both the DBC transaction and TEK state machine. When the first CM on a given downstream channel or bonding group joins a multicast session, the CMTS forwards the encrypted multicast traffic upon completion of the TEK state machine to ensure that the CM will not discard the encrypted multicast traffic. Alternatively, the CMTS may forward encrypted multicast traffic after sending the DBC-REQ but prior to receipt of the DBC-RSP from the CM. By not waiting for the DBC-RSP from the CM, the CMTS can start the encrypted multicast traffic sooner which will remove any delays induced by waiting for the DBC-RSP.

If the Security Association Encodings of a DSID changed and the CMTS does not receive the DBC-RSP after all the DBC-REQ retries and the CMTS has not received a TEK-request from the CM, the CMTS does not know whether the CM has implemented the DBC-REQ or not. The CMTS MUST recover from this condition. The CMTS actions may include the following for addition or deletion of a Security Association: initiation of a new DBC transaction to retry the errored DBC transaction, initiation of a new DBC transaction to undo the errored DBC transaction or reinitialization of the CM using the CM-CTRL-REQ message. Recovery is considered complete if the state of the Security Association is synchronized at both the CMTS and CM.

11.5.1.4 Changes to the Transmit Channel Set

Using DBC messaging, the CMTS can add channels to the transmit channel set, delete channels from the transmit channel set or replace one channel with another. Multiple actions can occur within a single DBC message. Whenever the CMTS changes the Transmit Channel Set, the CMTS MUST appropriately modify the SIDs associated with affected service flows. If the CM receives a DBC-REQ that causes a mismatch where one or more channels needed for a service flow are not included in the TCS, the CM MUST reject the DBC-REQ. For example, if service flow A is bonding across upstream channels 1, 2 and 3 and the CM receives a DBC-REQ to remove channel 1 and the DBC-REQ does not include the removal of SIDs associated with channel 1 for service flow A, the CM would reject the DBC-REQ message.

The CMTS MAY add channels to the TCS without specifying that these channels be used by any specific service flow. This allows the CMTS to add channels to the CM before they are actually needed to support service at that CM. If the CM receives a DBC-REQ that would result in more channels in the TCS than are needed to support the CM's service flows, the CM MUST NOT reject the DBC message due to the extra channel(s) unless the resulting TCS is inconsistent with the CM's transmit capabilities.

When the CMTS replaces a channel within the TCS, there are additional requirements beyond those of merely adding a channel combined with deleting a channel. These additional requirements exist because the service flow may be adversely affected during a channel replacement. From a process perspective, a channel replacement contains a channel deletion (the channel being replaced) and a channel addition (the replacement channel).

Clause 11.5.3 details the operation of the CMTS and CM during the DBC process. In the event of a corresponding SID Cluster change, the CMTS and the CM will follow the request-grant process detailed in clause 11.5.1.5. The CMTS then sends the DBC-REQ and sends ranging opportunities where applicable on any channels being added to the TCS. The CMTS then waits for the DBC-RSP. When the CM receives the DBC-REQ, it makes the transmit channel set changes identified in the DBC-REQ by immediately deleting any channels being removed from the TCS and applying the appropriate initialization technique to any channels being added to the TCS. Once the CM has successfully added a channel to its TCS, it begins using that channel for requesting and responding to grants if that channel is used by any of the CM's service flows. Once the CM has completed all attempts to add all new channels and deletes any channels being removed from the TCS, the CM sends a DBC-RSP. The DBC-RSP will contain no errors if the CM was able to make all TCS changes, partial service if the CM was able to make some of the TCS changes, failure if the CM was unable to make any TCS changes or rejection if the CM considers the DBC-REQ was invalid. Once the CMTS receives the DBC-RSP, it follows the process detailed in clause 11.5.1.5 before sending a DBC-ACK. The CMTS continues to accept requests and data transmissions received on deleted channels until the expiration of the T10 timer.

11.5.1.4.1 Impact of TCS Changes on Periodic Ranging

When the CMTS is removing a channel from a CM's TCS, the CMTS MUST continue sending station maintenance opportunities to the CM for the channel being removed until the CMTS receives the DBC-RSP from the CM. If the CMTS meets the maximum number of retries for invited ranging retries on the channel being removed during this period (DBC-REQ to DBC-RSP), the CMTS MUST NOT log this as an error condition because the CM may be in the process of removing this upstream channel. The purpose of the CMTS continuing to send the invited ranging opportunities is to ensure that the CM does not have a T4 expiration prior to processing the DBC-REQ message.

Similarly, when adding a new channel to the TCS with initialization technique of station maintenance or use directly, the CMTS MUST send the ranging opportunities while waiting for the DBC-RSP. These initialization techniques are used to shorten the DBC transaction time. Since these ranging opportunities can occur prior to the CM processing the DBC-REQ, the CMTS MUST NOT count these opportunities towards the Invited Ranging Retries (annex B) prior to receiving the DBC-RSP from the CM.

11.5.1.4.2 Exception Conditions for TCS Changes

When changing the TCS, error conditions may result in the CMTS never receiving the DBC-RSP. Recovering from this condition is up to CMTS vendor implementation. For example, the CMTS actions may include the following for a channel replacement or channel add:

- If the CMTS has sent RNG-RSP success on all new channels for the CM, the CMTS may assume DBC transaction success and assume the CM is operational on the new channels and has deleted the old channels.
- If the CMTS sends RNG-RSP success on only some of the new channels, the CMTS can assume that the CM was unable to acquire the remaining channels and is in the partial service mode of operation.
- If the CMTS sees the CM transmit on one or more new channels but the CM does not range successfully on any of the new channels, the CMTS knows the CM received the DBC-REQ and assumes the CM deleted the old channels, cannot use the new channels and is in partial service mode.
- If the CMTS does not see the CM transmit on any new channel, the CMTS assumes the CM never received the DBC-REQ. The CMTS can delete the new resources and reinstate the old resources.

The CMTS may attempt to remedy any partial service state by any combination of the following:

- Sending another DBC transaction to add missing channels.
- Forcing the CM to reinit MAC.
- Moving the CM to a different set of channels.

If the CM fails to receive a DBC-ACK after exhausting the retries for a DBC, the CM logs the error that the DBC-ACK was not received and proceeds to operate as if the DBC-ACK was received.

11.5.1.5 Changes to the Service Flow SID Cluster Assignments

Using the Service Flow SID Cluster Assignments TLV in the DBC messaging, the CMTS can assign new channels to a service flow, remove channels from a service flow or replace one channel with another for a service flow. Multiple actions can occur within a single DBC message.

Clause 11.5.3 details the operation of the CMTS and CM during the DBC process. Immediately after sending the DBC-REQ, the CMTS will start accepting bandwidth requests on new SIDs added by the Service Flow SID Cluster Assignment. If the overlap between the old and the new SID Clusters provides sufficient bandwidth as described in clause 11.5.1.5.1, the CMTS will stop granting on SIDs to be removed. If the overlap between the old and the new SID Clusters does not provide sufficient bandwidth, the CMTS will continue to grant bandwidth to the old SID Cluster. In either case, the CMTS will still accept bandwidth requests on SIDs to be removed from the old SID Cluster.

While waiting for the DBC-RSP, if the CMTS receives a bandwidth request using a SID that was newly added by the Service Flow SID Cluster Assignment or sends a RNG-RSP with confirmation code "success" on any new channel added in the TCS, then it will:

- Begin granting bandwidth to any SIDs added by the SF SID Cluster Assignment for channels which are ranging complete.
- Stop accepting requests from any SIDs deleted by the SF SID Cluster Assignments.
- Stop granting bandwidth to channels deleted from the TCS.
- Stop granting to any SIDs removed by the SF SID Cluster Assignment if there is sufficient bandwidth.

When the CM receives the DBC-REQ, it stops requesting on channels removed by the Service Flow SID Cluster Assignment, but continues to transmit data in any grants on these channels. The CM starts using any new channels for requesting, prepares to receive grants for these channels and sends a DBC-RSP. Once the CMTS receives the DBC-RSP with a confirmation code of okay or partial service, it will stop providing grants as well as accepting requests over the SIDs to be removed (if it has not done so already). Additionally, it will start providing grants using the new SIDs added by the Service Flow SID Cluster Assignment. The CMTS waits a vendor-specific time before sending the DBC-ACK to ensure that the CM is able to transmit in any grants outstanding for SIDs removed by the Service Flow SID Cluster Assignments. The CMTS then sends a DBC-ACK. When the CM receives the DBC-ACK, it removes the SIDs associated with any channels deleted by the Service Flow SID Cluster Assignment.

When the CMTS is not changing the TCS but is changing the Service Flow SID Cluster Assignment, error conditions may result in the CMTS never receiving the DBC-RSP. Recovering from this condition is up to CMTS vendor implementation. The CMTS actions may include the following for a Service Flow SID Cluster Assignment change:

- Attempting another DBC transaction.
- Forcing the CM to reinit the MAC.
- Initiating DSD messaging for the service flows possibly impacted.

If the CM fails to receive a DBC-ACK after exhausting the retries for a DBC transaction not changing the TCS, the CM logs the error that the DBC-RSP was not received. The CM MUST delete the SIDs for any Service Flow SID Cluster Assignment deletions. Thus, the CM stops responding to grants on any channels deleted by the Service Flow SID Cluster Assignment.

11.5.1.5.1 Bandwidth Sufficiency

When modifying the set of channels associated with a service flow, the CMTS determines whether or not there is sufficient bandwidth to adequately support the affected service flow during the DBC operation. The definition of sufficiency is left up to CMTS vendor implementation. Consider the following examples of a Service Flow SID Cluster Assignment change which replaces upstream channel 3 with upstream channel 9 for three service flows:

- Service flow B is bonded over upstream channels 1, 2 and 3. The CMTS looks at the quality of service parameters for service flow B and the bandwidth typically available on channels 1 and 2 to determine if there is sufficient bandwidth on these channels to adequately support the affected service flow. The CMTS sees that service flow B is best effort with no guaranteed minimum bandwidth and determines that there is sufficient bandwidth on channels 1 and 2 to meet the needs of this service flow during the DBC transaction. Hence, this would be a sufficient bandwidth case.
- Service flow C is also bonded across channels 1, 2 and 3. Service flow C has a guaranteed minimum bit rate of 5 Mbps. The CMTS determines that it needs to support this service during the DBC transaction and that there is insufficient bandwidth on channels 1 and 2 to meet the needs of this service flow. Thus, this would be an insufficient bandwidth case.
- Service flow D is a UGS service provisioned for channel 3. The CMTS determines that there is insufficient bandwidth to sustain the UGS flow during the DBC transaction because no service would be available between the time channel 3 is removed and channel 9 is actually added. Thus, this would be an insufficient bandwidth case.

This notion of sufficiency is a CMTS notion and impacts the Service Flow SID Cluster Assignment change process at the CMTS. Whenever the CMTS decides that there is insufficient bandwidth to adequately support a service flow during the replacement, the CMTS MAY send duplicate grants over the new and old channel sets during the DBC transaction. In the example of service flow D above, the CMTS would send grants on channel 9 and channel 3 during the DBC transaction to minimize the downtime of the service flow.

With TCS modifications, the CM deletes any old channels and adds any new channels upon receipt of the DBC-REQ. Receipt of the DBC-ACK for these cases serves only to stop the "DBC-ACK Timeout" timer.

11.5.1.6 Initialization Technique

There are many factors that drive the selection of an initialization technique when commanding a dynamic bonding change. While it is desirable to provide the minimum of interruption to existing QoS services such as voice over IP or streaming video sessions, a CM will not be able to successfully execute a channel change given an initialization technique that is unsuitable for the cable plant conditions. In some cases, a CM will impair the new channel given an unsuitable initialization technique. For instance, consider the use of initialization technique 4 (use the new channel(s) directly) when there is a significant difference in propagation delay between the old and new channels. Not only will the CM be unable to communicate with the CMTS on that channel after the channel change, but its transmissions may also interfere with the transmissions of other CMs using the channel.

Careful consideration needs to be given to the selection of an initialization technique. Some restrictions are listed below. This list is not exhaustive, but is intended to prevent the use of a particular initialization technique under conditions where its use could prevent the CM from successfully executing the channel change. Packets may be dropped under some conditions during channel changes; 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.

11.5.1.6.1 Initialization Technique One (1)

The use of initialization technique 1 (broadcast initial ranging) may result in a lengthy interruption of service. However, this interruption of service is mitigated by the reservation of QoS resources on the new channel(s). The service interruption can be further reduced if the CMTS supplies the UCD TLV in the DBC-REQ in addition to providing more frequent initial ranging opportunities on the new channel.

11.5.1.6.2 Initialization Technique Two (2)

The use of initialization technique 2 (unicast ranging) offers the possibility of only a slight interruption of service. In order to use this technique, the CMTS MUST include the UCD TLV in the DBC message if the upstream channel is changing.

However, the CMTS MUST NOT use this technique if:

- The DBC-REQ message contains an RCC that affected the CM's Primary Downstream Channel and that change results in a timing change.
- Propagation delay differences between the old and new channels will cause the CM burst timing to exceed the ranging accuracy requirements of [12] and the CMTS does not compensate for this difference with the ranging offset TLVs (clauses C.1.5.1.8.2 to C.1.5.1.8.5 on Ranging Offset TLVs).
- Attenuation or frequency response differences between the old and new upstream channels will cause the received power at the CMTS to be outside the limits of reliable reception and the CMTS does not compensate for this difference with the Power Offset TLVs (clause C.1.5.1.8.4 on Power Offset TLV).

11.5.1.6.3 Initialization Technique Three (3)

The use of initialization technique 3 (broadcast or unicast ranging) offers the possibility of only a slight interruption of service. This value might be used when there is uncertainty when the CM may execute the DBC command and thus a chance that it might miss station maintenance slots. However, the CMTS MUST NOT use this technique if the conditions for using techniques 1 and 2 are not completely satisfied.

11.5.1.6.4 Initialization Technique Four (4)

The use of initialization technique 4 (use the new channel directly) results in the least interruption of service.

In order to use this technique, the CMTS MUST:

- Synchronize timestamps and downstream symbol clocks across the Primary Downstream Channels involved.
- Include the UCD TLV in the DBC message if the upstream channel is changing.
- However, the CMTS MUST NOT use this technique if The modulation Rate rate changes whenif replacing one S-CDMA channel with another S-CDMA channel.
- Primary The primary Downstream downstream channel is being changed or affected by implicit or explicit changes in the Receive Module.
- The DBC-REQ message requires the CM to switch a channel or channels between S-CDMA and TDMA.
- Propagation delay differences between the old and new channels will cause the CM burst timing to exceed the ranging accuracy requirements of [12] and the CMTS does not compensate for this difference with the ranging offset TLVs (clauses C.1.5.1.8.2 to C.1.5.1.8.5 on Ranging Offset TLVs).
- Attenuation or frequency response differences between the old and new upstream channels will cause the received power at the CMTS to be outside the limits of reliable reception and the CMTS does not compensate for this difference with the Power Offset TLVs (clause C.1.5.1.8.4 on Power Offset TLV).
- Micro-reflections on the new upstream channel will result in an unacceptable PER (greater than 1 %) with the pre-equalizer coefficients initialized according to [12].

11.5.1.7 Fragmentation of DBC-REQ Messages

If the CMTS fragments the DBC-REQ message, it MUST ensure that the fragments arrive in order at the CM, as the CM is not required to resequence out-of-order DBC-REQ message fragments. The CMTS may do so either by sending all message fragments on a single downstream or by transmitting fragments such that individual channel latencies do not affect fragment order.

Upon receiving the first fragment of a DBC-REQ message, the CM starts a "DBC-REQ Timeout" timer. If the timer expires before all fragments of the DBC-REQ message have been correctly received, the CM sends a DBC-RSP with confirmation code error-DBC-REQ-incomplete, then returns to the operational state. Correct reception of the DBC-REQ message fragments could include in-order reception of all fragments.

11.5.2 Exception Conditions

The CM MUST reject a message that the CM determines to be invalid or inconsistent with the CM's capabilities and service flows.

A DBC-REQ is considered invalid if any of the following apply:

- The message format does not match the format required for a DBC-REQ.
- The DBC-REQ contains an RCC that affects the CM's primary downstream but does not contain an initialization technique.
- The DBC-REQ includes an RCS change but does not specify one and only one downstream channel to be the Primary Downstream Channel.
- A CMTS-initiated DSA, DSC or DCC transaction is in progress at the CM.
- The DBC-REQ contains TCC Encodings which cause a violation of the Dynamic Range Window.
- The DBC-REQ contains a TCC Encoding with an Upstream Channel Action of Add (1) or Replace (4), but does not contain a UCD.

A DBC-REQ is considered inconsistent with the CM's capabilities and service flows if any of the following apply:

- Implementation of the DBC-REQ would require more downstream receivers than the CM has available.
- Implementation of the DBC-REQ would require the CM to switch from legacy mode to Multiple Transmit Mode.
- Implementation of the DBC-REQ would require more upstream transmitters than the CM has available.
- Implementation of the tuning range required by the DBC-REQ is inconsistent with the CM's capabilities.
- Implementation of the DBC-REQ would require different physical-layer implementation than the CM has available.
- Implementation of the DBC-REQ would require more SID Clusters than the CM supports.
- Implementation of the DBC-REQ would require more DS Resequencing DSIDs than the CM supports.
- Implementation of the DBC-REQ would require more DSIDs than the CM supports.
- Implementation of the DBC-REQ would require an RCS change that is inconsistent with the DS Resequencing Channel List.
- Implementation of the DBC-REQ would require a DS Resequencing Channel List change that is inconsistent with the RCS.
- Implementation of the DBC-REQ would require a TCS change that is inconsistent with the Service Flow SID Cluster Assignment.
- Implementation of the DBC-REQ would require a Service Flow SID Cluster Assignment that is inconsistent with the TCS.

- Implementation of the DBC-REQ would require more DSIDs with multicast attributes than the CM supports.
- The DBC-REQ message contains a client MAC address that is not in the CM's forwarding table.

If the CM considers the DBC-REQ message to be valid but is unable to acquire new downstream channels in the RCS and/or new upstream channels in the TCS, the CM responds with a DBC-RSP <Partial Service>.

If the CM is unable to acquire one or more downstream channels, the CM sends a DBC-RSP <Partial Service> and enters a partial service mode of operation in the downstream (see clause 8.4). Likewise, if the CM is unable to acquire one or more upstream channels, the CM sends a DBC-RSP <Partial Service> and enters a partial service mode of operation in the upstream (see clause 8.4).

If a CM issues a DSA-REQ or DSC-REQ for more resources and the CMTS needs to do a DBC 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-DBC" (refer to clause C.4) in the DSA-RSP or DSC-RSP message to indicate that the new resources will not be available until a DBC is received. The CMTS will then follow the DSA or DSC transaction (expiration of T10 transaction timer) with a DBC transaction.

The CMTS MUST NOT issue a DBC command to a CM if a DSA, DSC or DCC transaction is still outstanding at that CM. The CMTS MUST NOT issue a DSA, DSC or DCC command to a CM if the CMTS has previously issued a DBC command to that CM and that command is still outstanding.

If the CMTS issues a DBC-REQ command to a CM and that CM simultaneously issues a DSA-REQ or DSC-REQ then the CMTS command takes priority. The CMTS MUST respond with a confirmation code of "reject-temporary" for the DSA-REQ or DSC-REQ, per annex B. If the CM receives a DBC-REQ prior to receiving a DSA-RSP or DSC-RSP, the CM assumes that the CMTS will reject the DSA or DSC transaction and the CM MUST execute the DBC command.

If the CMTS sends a DBC-REQ and does not receive a DBC-RSP prior to the expiration of the Initializing Channel Timeout, it MUST retransmit the DBC-REQ up to a maximum of "DBC-REQ Retries" (annex B) before declaring the transaction a failure. Note that if the DBC-RSP was lost in transit and the CMTS retries the DBC-REQ, the CM may have already changed channels.

If the CMTS receives a DBC-RSP with confirmation code "error-DBC-REQ-incomplete", it determines whether "DBC-REQ Retries" has been exhausted before resending the DBC-REQ or declaring the transaction a failure.

If the CM sends a DBC-RSP and does not receive a DBC-ACK from the CMTS within, "DBC-ACK Timeout", it MUST retry the DBC-RSP up to a maximum of "DBC-RSP Retries" (annex B).

The CM MUST consider the DBC-REQ as a redundant command if the CM receives a DBC-REQ with any of the following:

- CM Receive Channel Configuration Encodings equal to the current Receive Channel Configuration.
- Any DSID encoding that adds an existing DSID.
- Transmit Channel Configuration Encoding that adds an upstream channel that is already present in the CM's Transmit Channel Set.

If the CM considers the DBC-REQ to be redundant, the CM MUST NOT execute the DBC-REQ. Then the CM MUST return a DBC-RSP, with a detailed confirmation code of "reject-already-there" to the CMTS per annex C.

If the CM does not receive a DBC-ACK after all the retries, the CM logs an error and continues normal operation.

11.5.3 DBC State Transition Diagrams

11.5.3.1 CMTS DBC State Transition Diagrams

The CMTS MUST support the DBC operation as shown in the State Transition diagrams in figure 11-50 through figure 11-53.

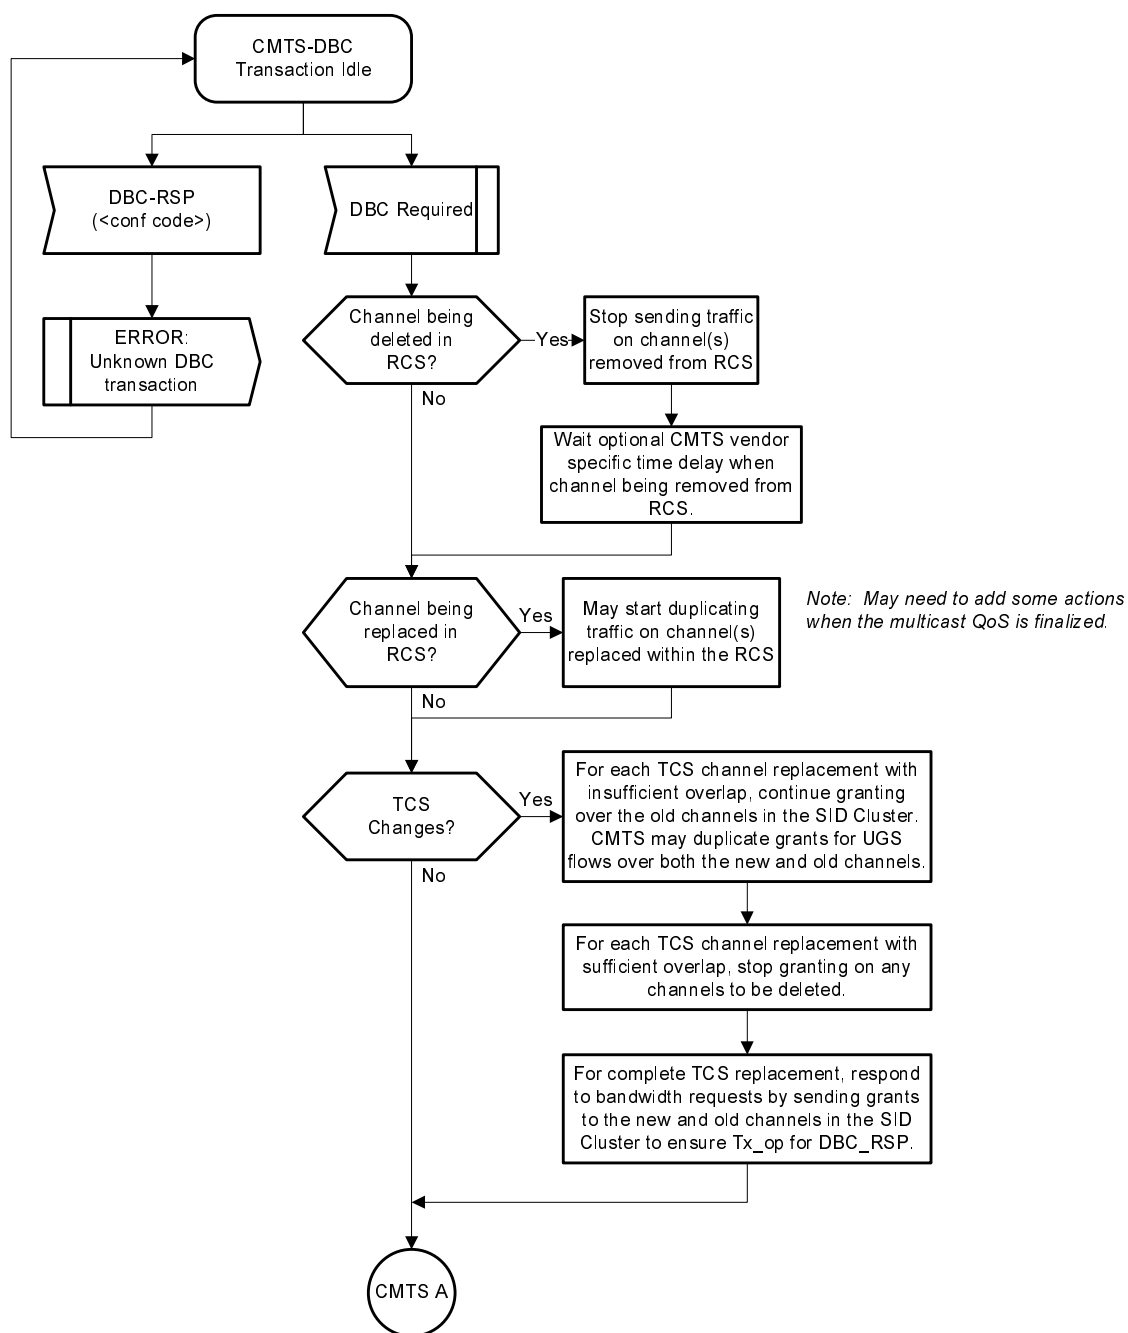


Figure 11-50: CMTS DBC Request (part 1)

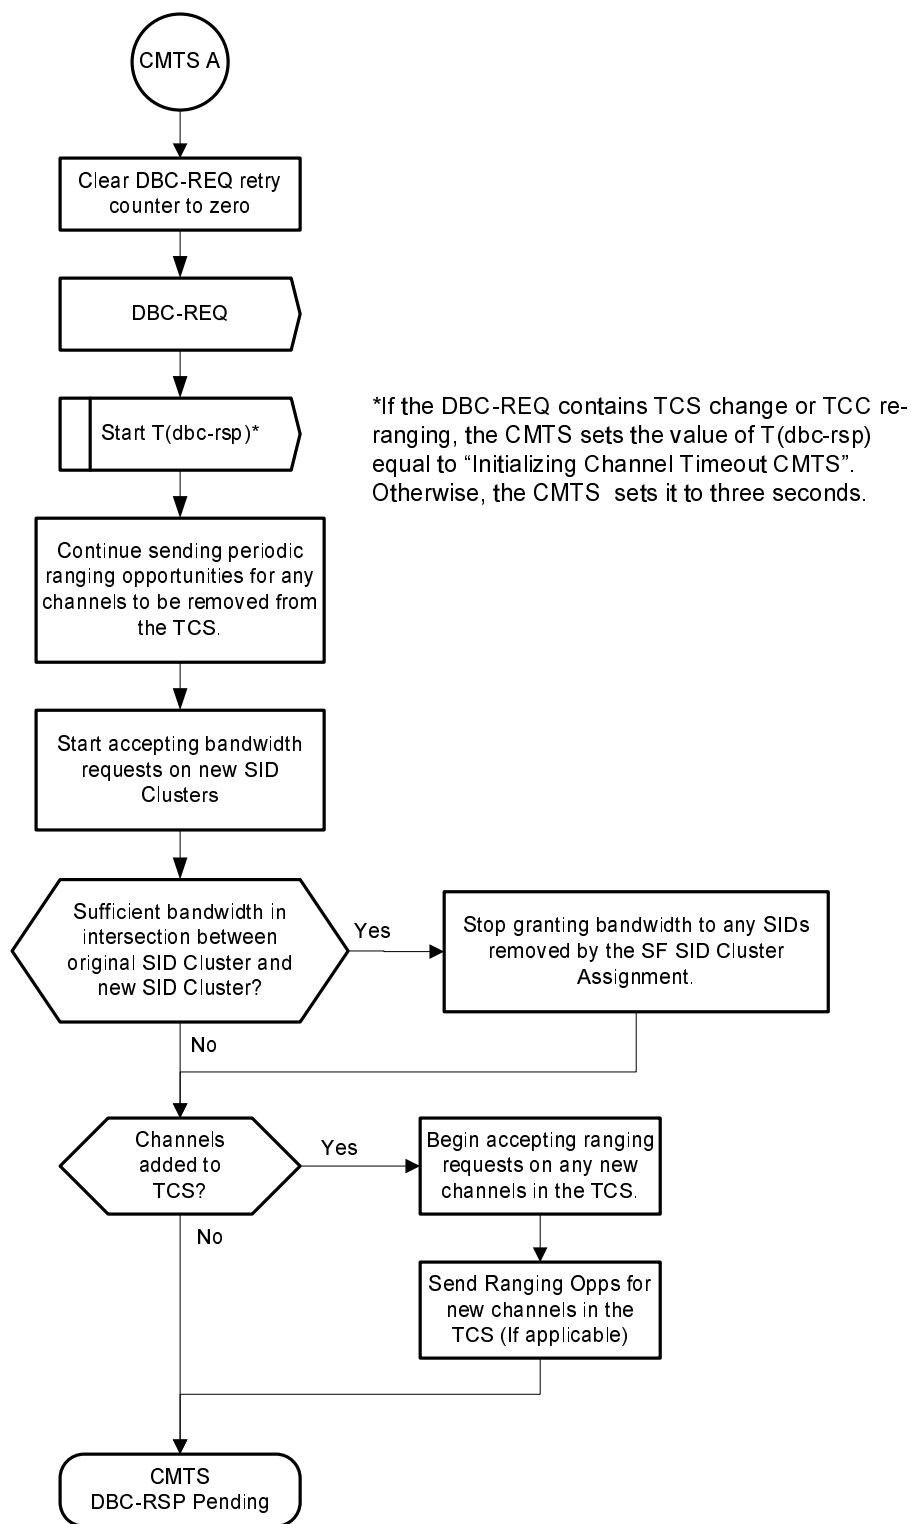


Figure 11-51: CMTS DBC Request (part 2)

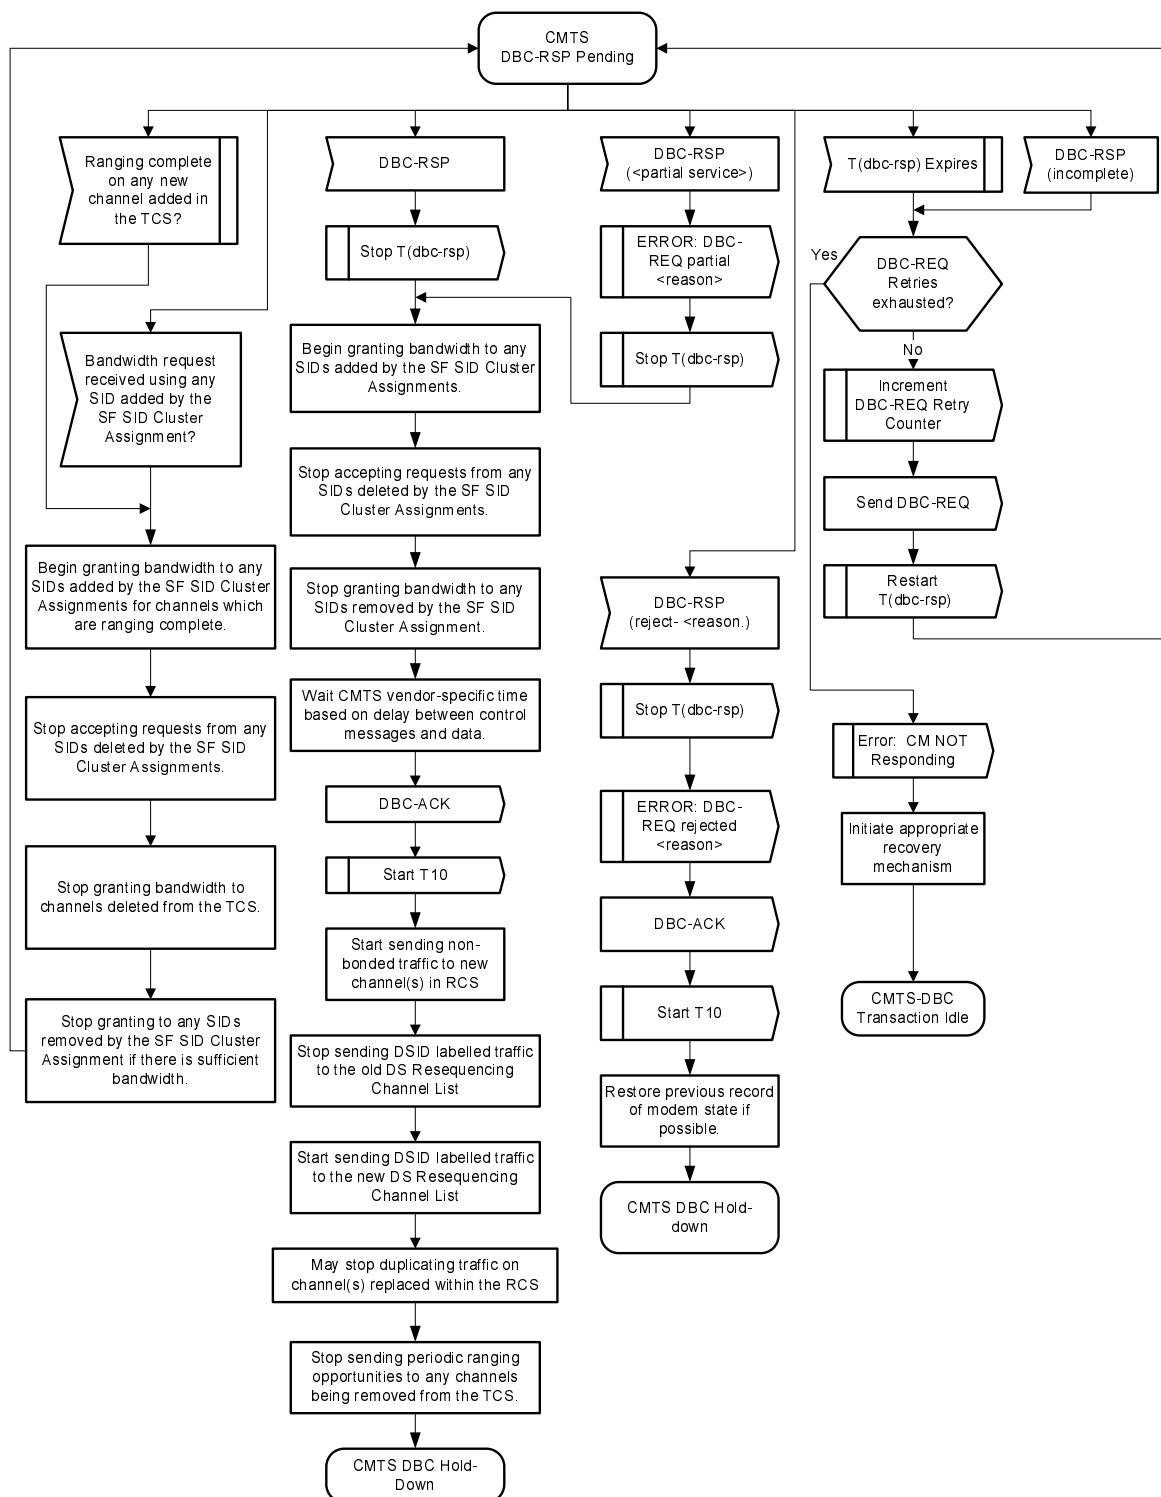


Figure 11-52: CMTS DBC-RSP Pending

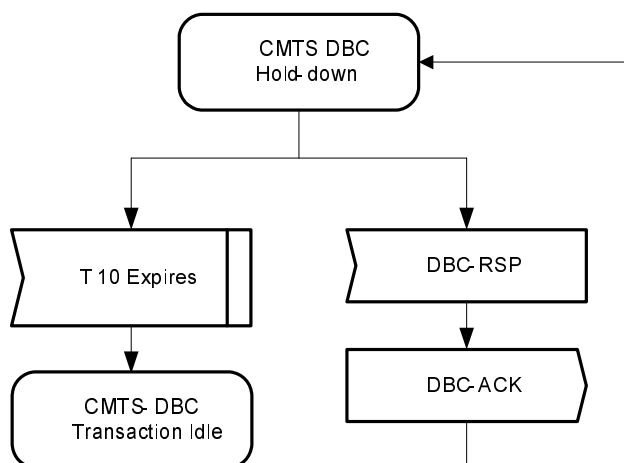


Figure 11-53: CMTS DBC Hold-down

11.5.3.2 CM DBC State Transition Diagrams

The CM MUST support the DBC operation as shown in the State Transition diagrams in figure 11-54 through figure 11-59.

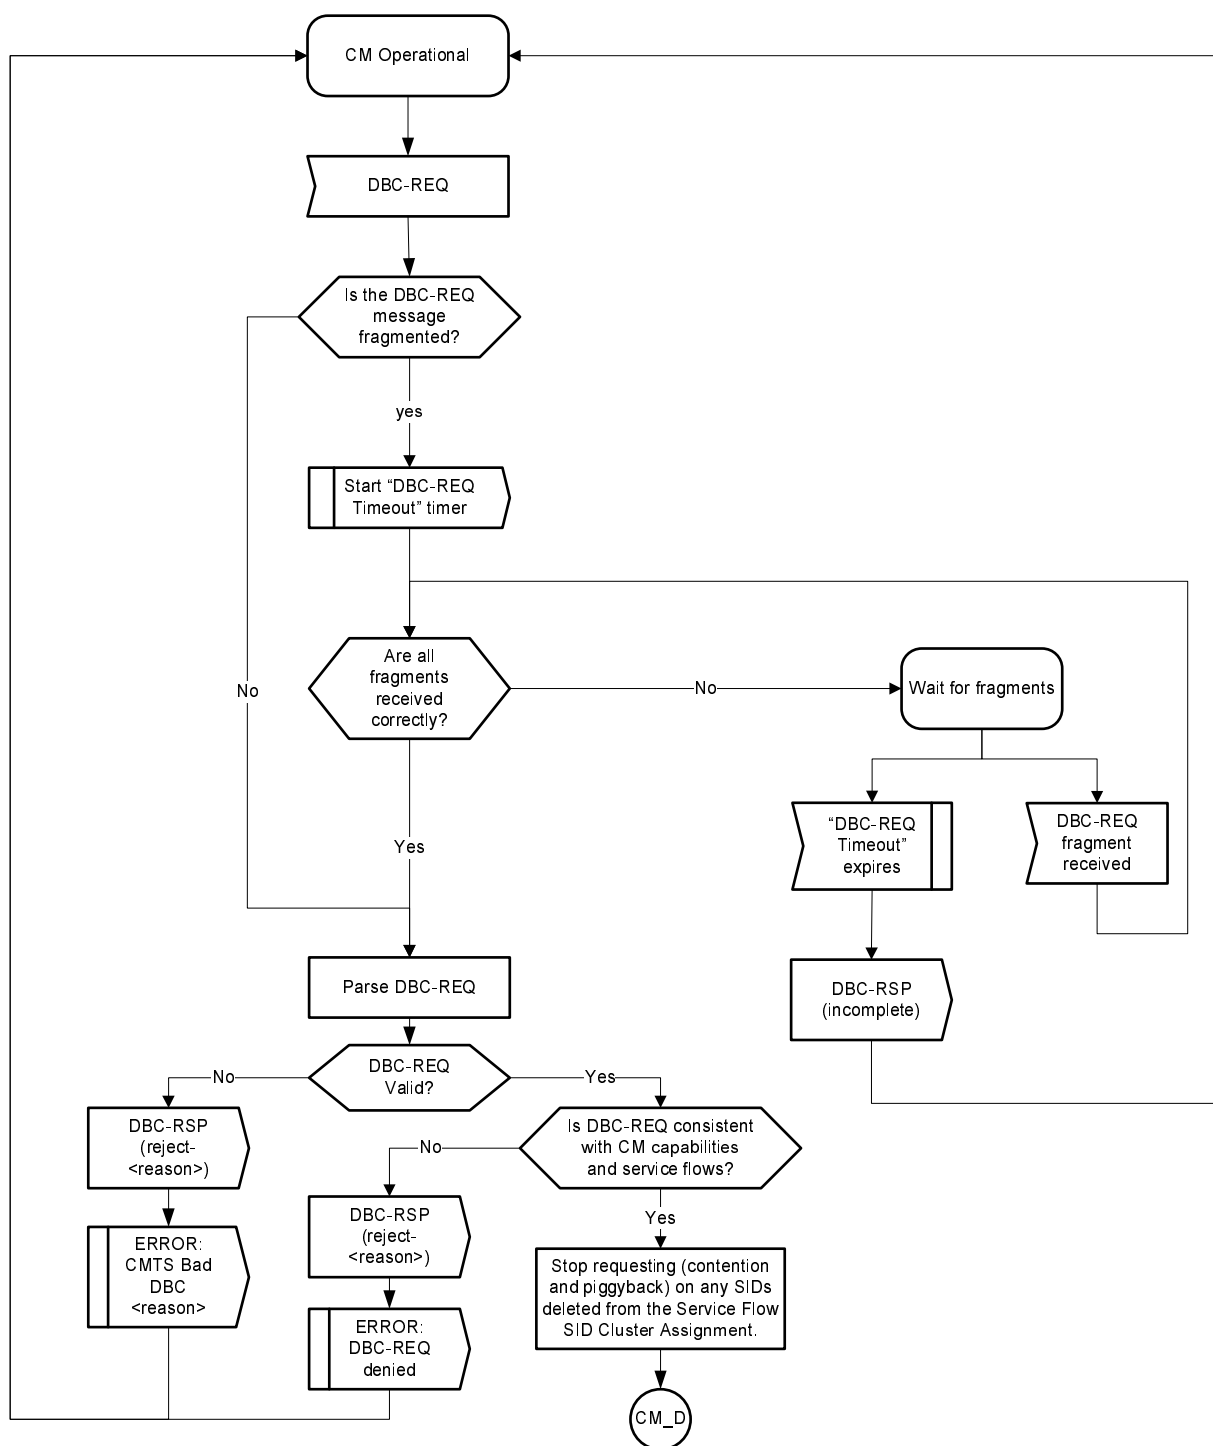


Figure 11-54: CM DBC-RSP (part 1)

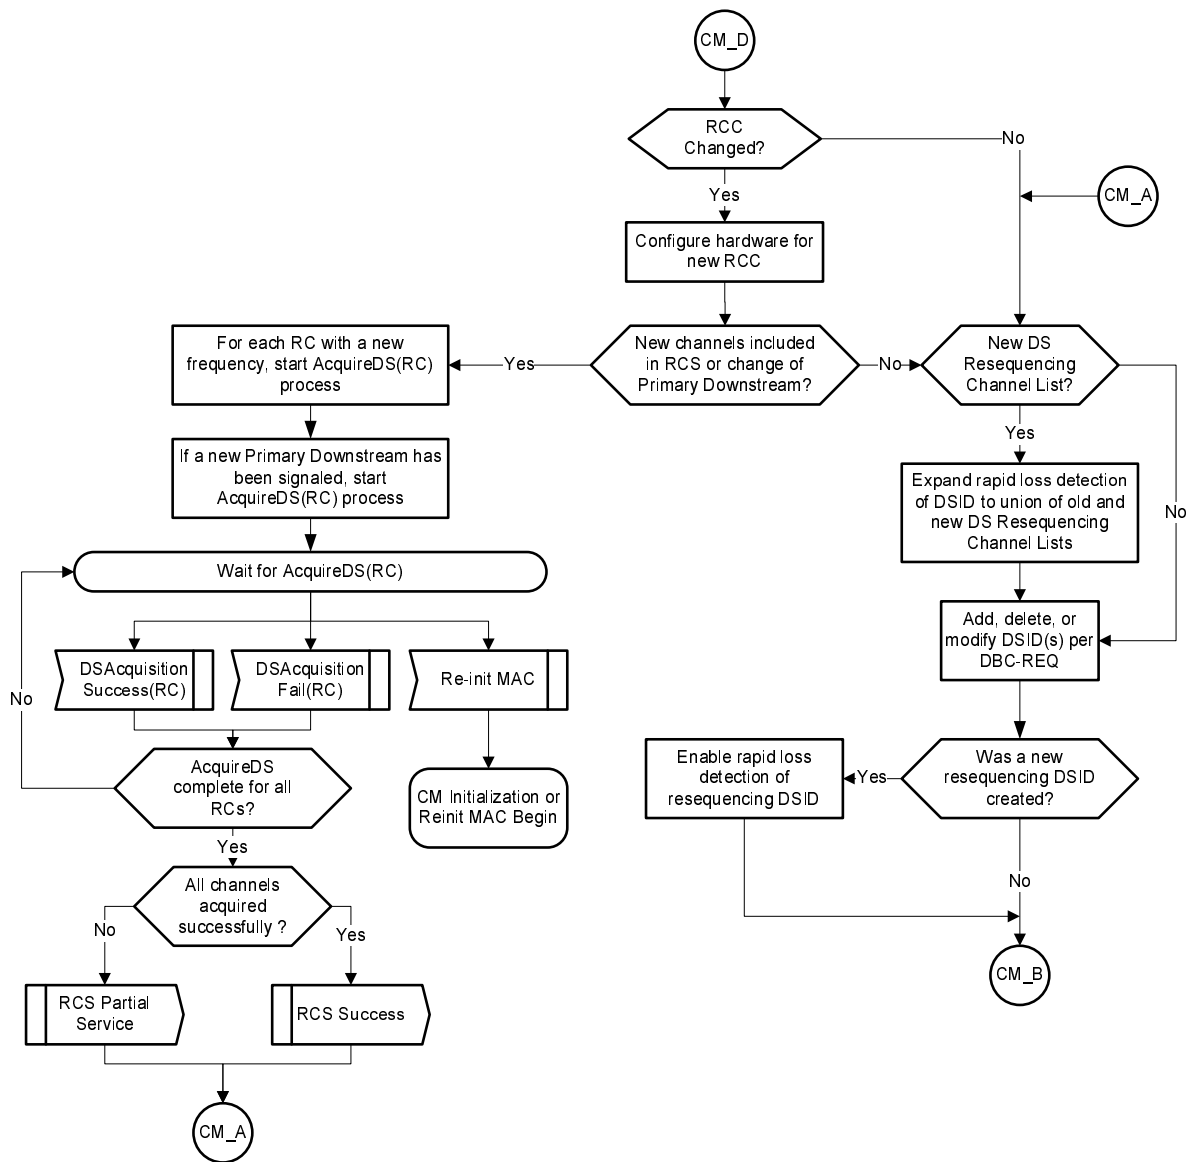
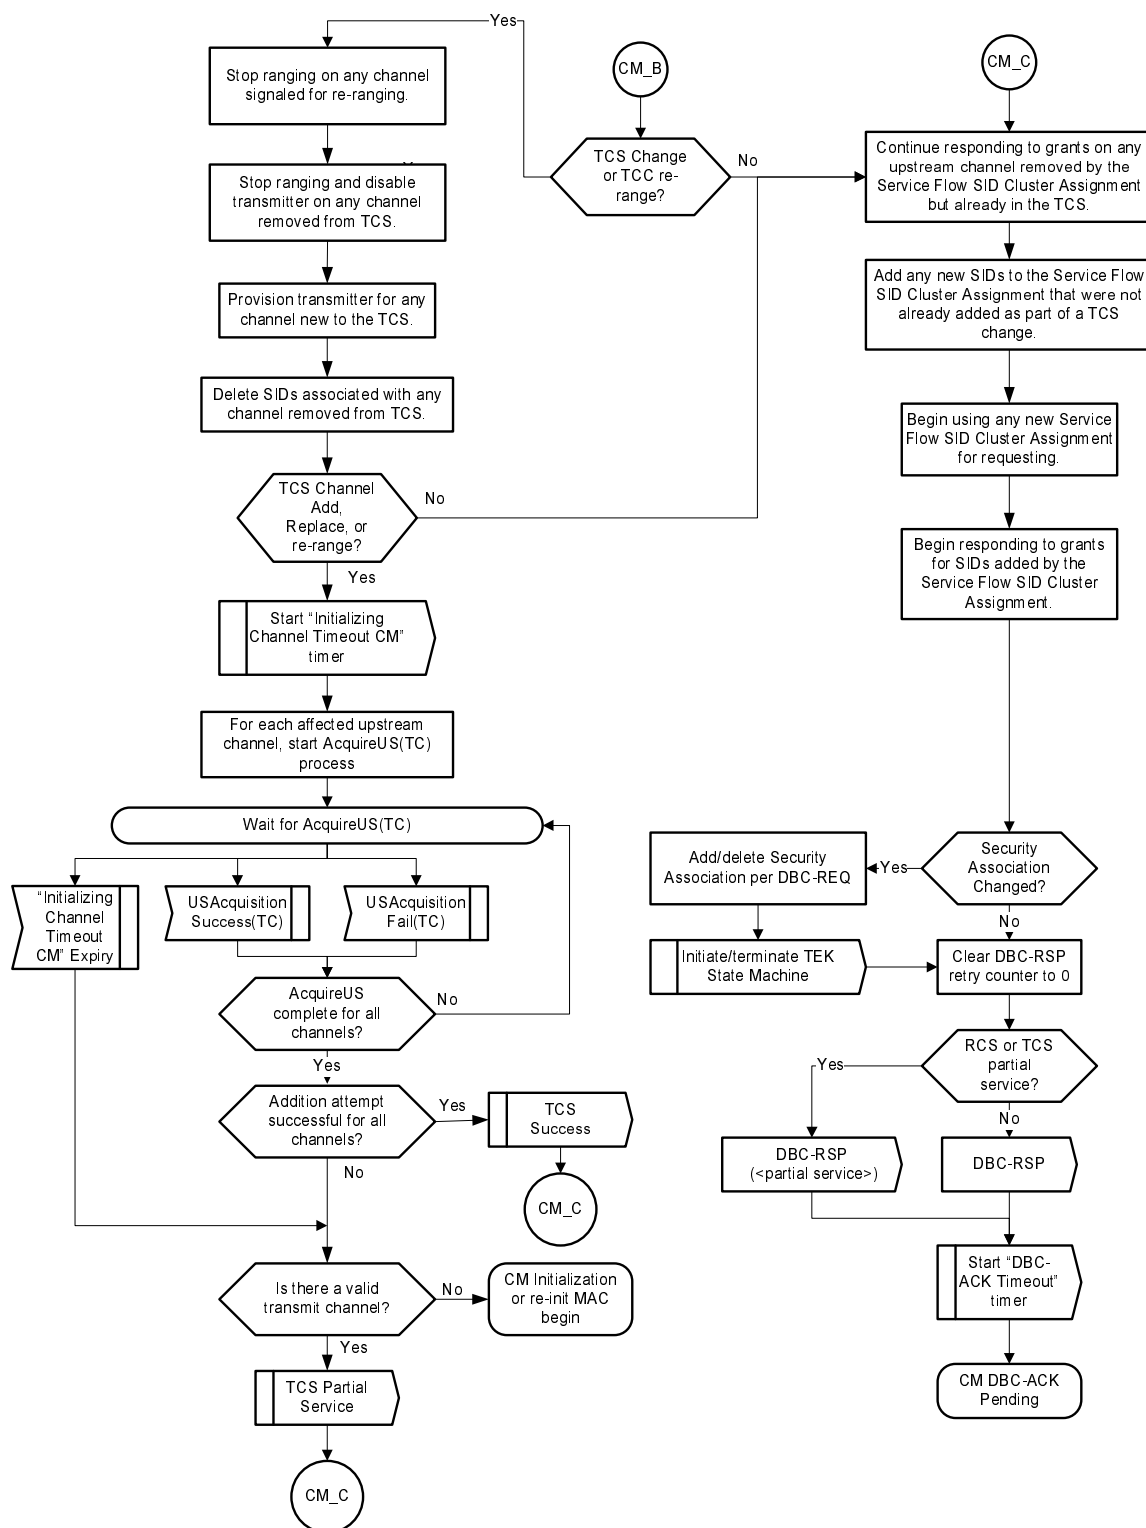
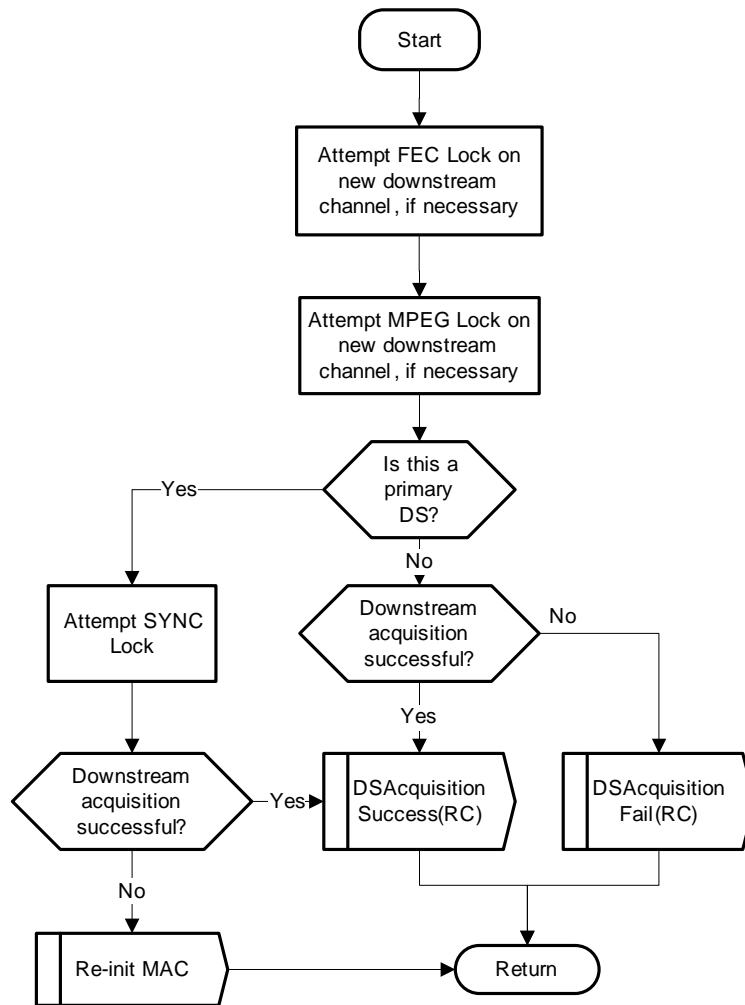
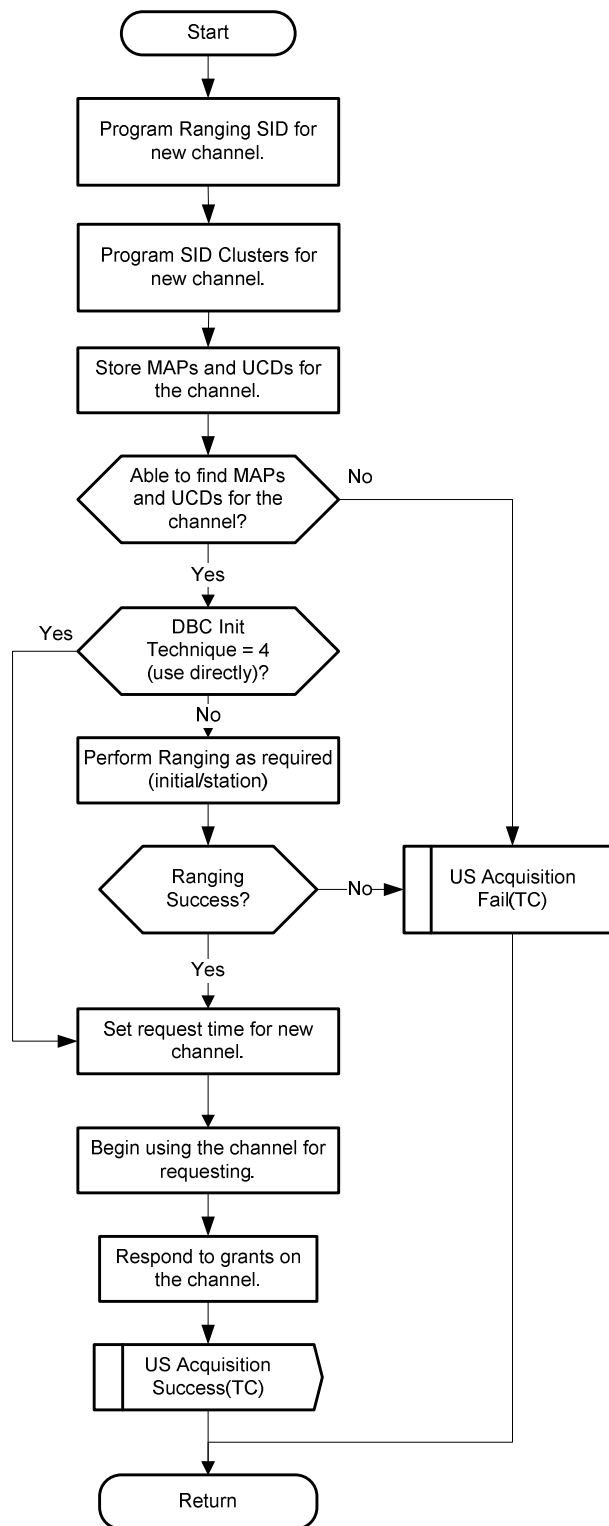


Figure 11-55: CM DBC-RSP (part 2)



**Figure 11-57: CM AcquireDS Procedure**

**Figure 11-58: CM AcquireUS Procedure**

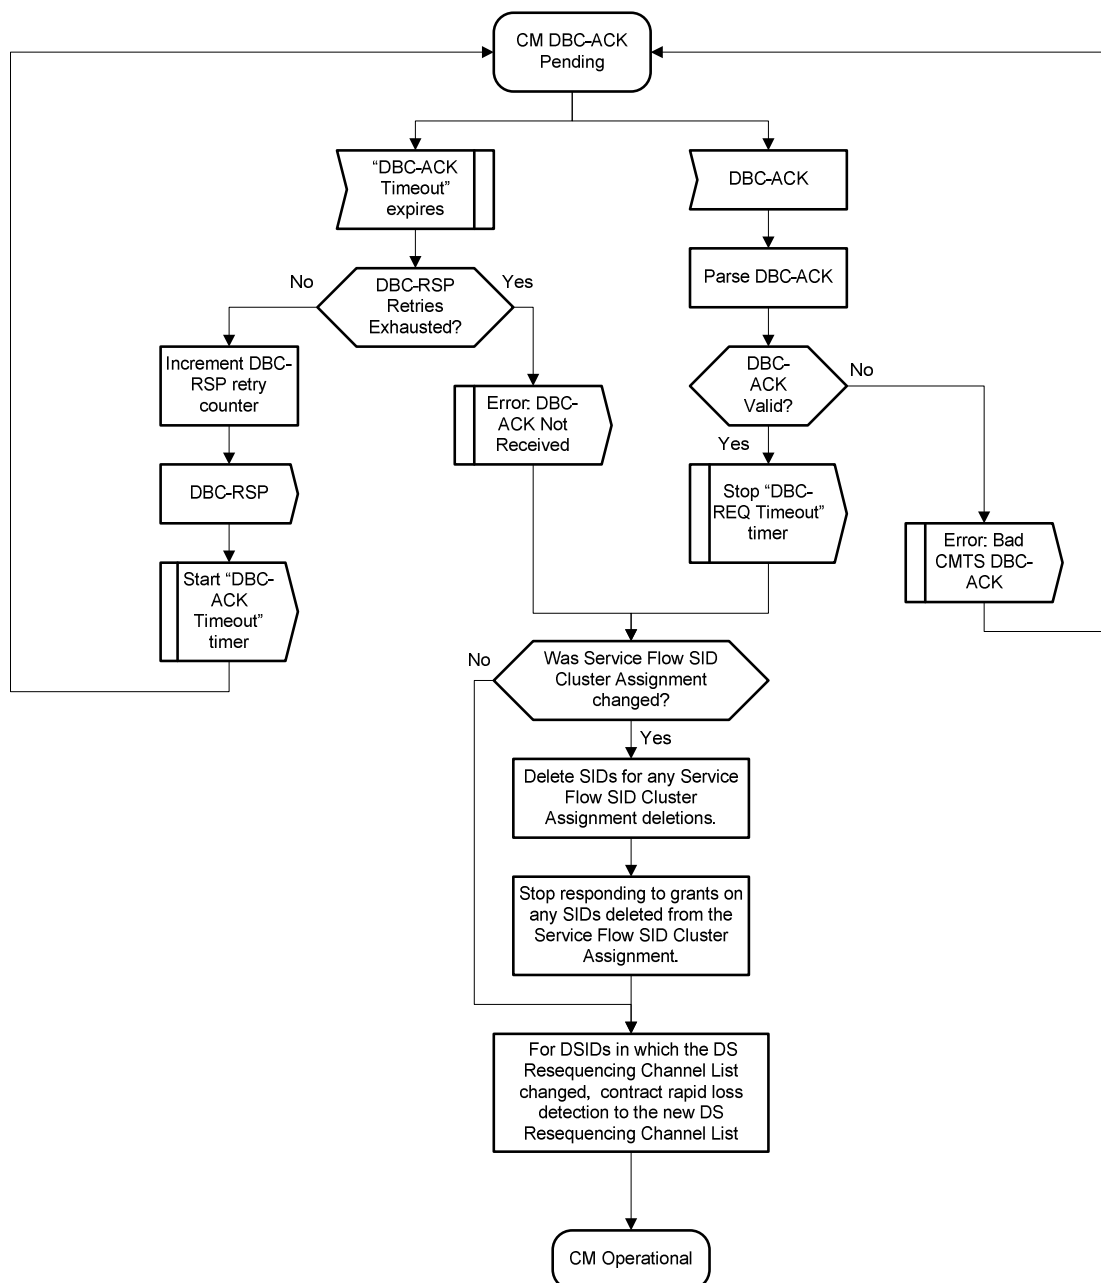


Figure 11-59: CM DBC-ACK Pending

11.6 Autonomous Load Balancing

Autonomous Load Balancing is a feature of the CMTS that controls dynamic changes to the set of downstream and upstream channels used by a CM.

The CMTS uses the Dynamic Channel Change (DCC) message to control the load balancing of CMs not operating in Multiple Receive Channel mode. The CMTS can also use DCC to load balance the upstream of a CM to which a Transmit Channel Configuration was not assigned in the registration process. The CMTS uses the Dynamic Bonding Change (DBC) message to control load balancing of CMs operating in Multiple Receive Channel mode. With CMs operating in Multiple Receive Channel mode, load balancing can be performed by changing the Receive Channel Set of the CM or by moving one or more service flows to different downstream channels within the current RCS of the CM. With CMs operating in Multiple Transmit Channel mode, load balancing can be performed by changing the Transmit Channel Set of the CM or by moving one or more service flows to different upstream channels within the current TCS of the CM.

11.6.1 Load Balancing Groups

A "Load Balancing Group" (LBG) is a set of upstream and downstream channels over which a CMTS performs load balancing for a set of CMs.

A Load Balancing Group has the following attributes:

- a set of downstream and upstream channels in the same CM Service Group (CM-SG);
- a policy which governs if and when a CM or its individual service flows can be moved; and
- a priority value which can be used by the CMTS in order to select which CMs to move.

The CMTS creates a Load Balancing Group for every MD-CM-SG that is instantiated by the topology configuration. This type of LBG is referred to as a "General" Load Balancing Group. Further the operator can configure "Restricted" Load Balancing Groups that contain a subset of the channels in a CM-SG to which a CM can be assigned.

The CMTS MUST support configuration of each channel (upstream and downstream) to more than one LBG (Restricted or General). The CMTS attempts to balance load among all of the channels of each LBG. In cases where a single channel (upstream or downstream) is associated with more than one LBG, the CMTS might have to consider all LBGs for which such overlaps exist in its load balancing algorithm.

During Registration, the CMTS attempts to assign each CM to a Load Balancing Group. If the operator has configured a CM to be in a Restricted Load Balancing Group, then the CMTS restricts the CM to the channels of the configured Restricted Load Balancing Group. If the operator has not configured a CM to be in a Restricted Load Balancing Group, but the CMTS can determine a General Load Balancing Group (i.e. MD-CM-SG) for the CM, the CMTS performs load balancing for the CM among the channels of the General Load Balancing Group. If the CMTS cannot determine either an assigned Restricted Load Balancing Group or the General Load Balancing Group for a registered CM, the CMTS does not perform autonomous load balancing of the CM.

The CMTS MUST NOT assign a CM to more than one Load Balancing Group.

11.6.1.1 General Load Balancing Groups

The CMTS MUST implement a General Load Balancing group for every MD-CM-SG, containing all channels of that MD-CM-SG.

The CMTS attempts to identify the General Load Balancing Group (MD-CM-SG) for a CM during the topology resolution process (see clause 10.2.3).

Every CM registers into an MD-CM-SG. The DOCSIS 3.0 initialization procedure enables a CMTS to automatically determine the MD-CM-SG of a DOCSIS 3.0 CM when it initially ranges. In many plant topologies, an upstream channel is configured into a single MD-CM-SG, so the CMTS can also determine the MD-CM-SG for a pre-3.0 DOCSIS CM from its single upstream channel. The CMTS MUST support load balancing of pre-3.0 DOCSIS CMs when their General Load Balancing Group (MD-CM-SG) is determined from the upstream/downstream channel pair upon which they range and register. The CMTS MAY support automatically determining the MD-CM-SG of a pre-3.0 DOCSIS CM when it is MD-CM-SG cannot be determined from the upstream/downstream channel pair. In the case where the MD-CM-SG cannot be determined, the CM is not associated with a General Load Balancing Group and is thus not eligible to be moved during load balancing operations unless it is assigned to a Restricted Load Balancing Group.

As explained in clause 5, the set of downstream channels within a MAC Domain that reach a single CM is called an MD-DS-SG. Similarly, the set of upstream channels within a MAC Domain that reach a single CM is called an MD-US-SG.

The default load balancing policy, available initialization techniques and enable/disable control for a General Load Balancing Group are configured for the GLBG's MAC Domain on the Fibre Node that GLBG serves. In most cases, a GLBG will serve a single Fibre Node, so this MAC Domain-Fibre Node pair maps to a single GLBG. If the GLBG serves multiple Fibre Nodes, the CMTS enforces that all are configured with the same default policy, initialization techniques and enable/disable control.

11.6.1.2 Restricted Load Balancing Groups

Restricted Load Balancing Groups are used to accommodate a topology specific or provisioning specific restriction (such as a set of channels reserved for business customers). The CMTS can associate an upstream or downstream channel with any number of Restricted Load Balancing Groups. A CM can be configured to an identified Restricted Load Balancing Group with the Service Type Identifier or the Load Balancing Group Identifier encodings in the CM configuration file (see clauses C.1.1.18.1.10 and C.1.1.18.1.3).

The CMTS **MUST** enforce that all Restricted LBGs are configured with channels within the same CM Service Group (CM-SG) (see clause 5.2.6). The CMTS **SHOULD** enforce that all Restricted LBGs are configured with channels within the same MAC Domain CM Service Group (MD-CM-SG). Load Balancing across MAC Domains is out of scope of the present document. The CMTS **MUST** enforce that a configured LBG contain both downstream and upstream channels. The CMTS **MUST** permit configuration of the channels of a Restricted Load Balancing Group to consist of some or all of the channels in an MD-CM-SG.

The CMTS will assign a modem to a Restricted Load Balancing Group only if it is explicitly provisioned (via CMTS management objects or configuration file TLV) to be a member of that group.

When the CMTS receives a Registration Request message, the CMTS **MUST** identify whether this CM has been configured to a specific Service Type or to a Restricted Load Balancing Group via CMTS management objects (see [10]). If the CM is not assigned to a Service Type or Restricted Load Balancing Group via CMTS management objects, the CMTS **MUST** check for the presence of the Service Type Identifier and CM Load Balancing Group TLVs in the Registration Request message to identify assignment to a Restricted Load Balancing Group. If the CM is assigned to a Service Type or Restricted Load Balancing Group via CMTS management objects, the CMTS **MUST** ignore both the Service Type Identifier and the CM Load Balancing Group TLV in the Registration Request message. If the Registration Request contains a Load Balancing Group ID that is not defined on the CMTS, the CMTS **MUST** ignore the group ID.

If the CM has been assigned to a Restricted Load Balancing Group (either via CMTS management objects or via the config file) and the CMTS detects that the CM is registering on a channel pair that is not associated with the assigned Load Balancing Group, the CMTS **MUST** move the CM to an appropriate set of channels in the assigned group (either via the channel assignment in the REG-RSP-MP or by initiating a DCC-REQ when registration completes).

11.6.2 CMTS Load Balancing Operation

When load balancing is enabled for a particular CM, the CMTS adheres to the following restrictions:

- If the CM is assigned to a Load Balancing Group, the CMTS **MUST NOT** direct the CM or any of its service flows to move to a channel outside the Load Balancing Group to which it is assigned.
- The CMTS **MUST** move the CM or its service flows to channels on which the CM can operate. The CMTS **MUST NOT** move a DOCSIS 1.0 or DOCSIS 1.1 compliant CM or a DOCSIS 2.0 compliant CM that has 2.0 mode disabled or a 3.0 CM with MTC Mode disabled and 2.0 mode disabled, to a Type 3 or Type 4 upstream channel. The CMTS **MUST** perform autonomous load balancing of CMs not operating in Multiple Receive Channel mode with a message supported by the CM, i.e. DCC-REQ (DOCSIS 1.1/2.0) or UCC-REQ (DOCSIS 1.0). The CMTS **MUST** be capable of performing intra-MAC Domain load balancing of CMs, operating in Multiple Receive Channel mode, either for the entire CM or any individual service flows of the CMs with a DBC message.
- As described in clause 8.1.1, the CMTS **MUST** ensure that the Required and Forbidden Attributes are met when moving the CM or its service flows.
- If the CMTS cannot determine a Load Balancing Group of the CM, the CMTS **MUST NOT** perform autonomous load balancing of the CM or any of its service flows.

The CMTS has many factors to consider when autonomously load balancing; these include primary downstream capability (and thus the availability of a DS channel for pre-3.0 DOCSIS CMs), MAP/UCD assignment to DS channels, attribute-based channel assignment, restricted load balancing groups and multicast replication requirements. When a CM is assigned to a restricted load balancing group, the CMTS **MUST** give that assignment precedence over the Service Flow attribute-based channel assignment and the multicast replication requirements.

If load balancing is disabled for a CM (either system-wide or for the load balancing group to which the CM is assigned or via the load balancing policy assigned to the CM) the CMTS adheres to the following restrictions:

- The CMTS MUST NOT perform autonomous load balancing of the CM.
- If the CM supports Multiple Receive Channel mode, the CMTS MUST assign (in Registration Response) an RCC for which the primary DS channel is the channel upon which the Registration Response is transmitted. If a suitable RCC cannot be provided, the CMTS MUST disable Multiple Receive Channel Mode.
- If the CM supports Multiple Transmit Channel mode, the CMTS MUST assign (in Registration Response) a Transmit Channel Set containing the upstream channel upon which the CM transmitted its Registration Request.

11.6.3 Multiple Channel Load Balancing

Operating in Multiple Transmit Channel and/or Multiple Receive Channel mode provides a level of load-balancing on its own. However, in cases where the number of downstream channels or upstream channels in the MD-CM-SG exceeds the number of receive channels or transmit channels for a particular CM, the CMTS performs load balancing using the Dynamic Bonding Change message.

For a CM operating in Multiple Transmit Channel mode, the CMTS performs autonomous load balancing by transmitting DBC messages that change the Transmit Channel Set of the CM and/or the SID_clusters of the CM's upstream service flows.

For a CM operating in Multiple Receive Channel mode, the CMTS performs autonomous load balancing by transmitting DBC messages that change the Receive Channel Configuration, DSIDs and/or Resequencing Channel Lists of the CM.

For a CM operating in Multiple Receive Channel mode, the CMTS can perform autonomous load balancing of a non-bonded, non-resequenced individual downstream service flow to a different downstream channel in the CM's Receive Channel Set without notifying the CM with a DBC message.

11.6.4 Initialization Techniques

The description of a Load Balancing Group includes the initialization technique(s) that could be used when autonomously load balancing a cable modem within the group. The initialization technique definition for each Load Balancing Group is represented in the form of a bit map, with each bit representing a specific technique (bits 0-4). Initialization technique 0 is only defined for DCC (not DBC). If a Load Balancing Group is restricted to only use initialization technique 0, the CMTS will be forced to use DCC for any CM that it attempts to move.

11.6.5 Load Balancing Policies

Load balancing policies allow control over the behavior of the autonomous load balancing process on a per-CM basis. A load balancing policy is described by a set of conditions (rules) that govern the autonomous load balancing process for the CM. When a load balancing policy is defined by multiple rules, all of the rules apply in combination. This specification does not intend to place requirements on the specific algorithms used by the CMTS for load-balancing, nor does it make a statement regarding the definition of "balanced" load. CMTS vendors are free to develop appropriate algorithms in order to meet market and deployment needs.

Load balancing rules and the load balancing policy definition mechanism have been created to allow for specific vendor-defined load balancing actions. However, there are two basic rules that the CMTS is required to implement.

The CMTS MUST implement the following basic rules:

- Prohibit load balancing using a particular CM.
- Prohibit load balancing using a particular CM during certain times of the day.

The policy ID value of zero is reserved to indicate the CMTS's basic load balancing mechanism, which does not need to be defined by a set of rules.

Each Load Balancing Group has a default load balancing policy. During the registration process, the CMTS MUST assign the CM a load balancing policy ID. The policy ID may be assigned to a cable modem via the cable modem config file. The CMTS MUST assign the CM the load balancing policy ID provisioned in the config file and sent in the Registration Request, if it exists. Otherwise, the CMTS MUST assign the CM the default policy ID defined for the Load Balancing Group.

The per-CM load balancing policy ID assignment can be modified at any time while the CM is in the operational state via internal CMTS processes and potentially via CMTS management objects; however, the policy ID is always overwritten upon receipt of a Registration Request message.

11.6.6 Load Balancing Priorities

A Load Balancing priority is an index that defines a rank or level of importance, which is used to apply differential treatment between CMs in the CMTS's load balancing decision process.

In general, a lower load balancing priority indicates a higher likelihood that a CM will be moved due to load balancing operations. The CMTS MAY take many factors into account when selecting a CM to move, of which priority is only one. When other factors are equal, the CMTS SHOULD preferentially move a CM with lower load balancing priority over one with higher load balancing priority.

The CMTS MUST associate each cable modem with a load balancing priority. Priority may be assigned to a cable modem via the cable modem config file. The CMTS MUST assign the CM the load balancing priority provisioned in the config file and sent in the Registration Request, if it exists. If a cable modem has not been assigned a priority, it is associated with the default (lowest) load balancing priority value of zero.

The per-CM load balancing priority assignment can be modified at any time while the CM is in the operational state via internal CMTS processes as dictated by a specific load balancing policy; or potentially via CMTS management objects; however, the priority assignment is always overwritten upon receipt of a Registration Request message.

11.6.7 Load Balancing and Multicast

In order to efficiently manage multicast traffic and balance load across a Load Balancing Group, it is reasonable to expect that the CMTS might attempt to reduce the amount of duplicated multicast traffic by consolidating all members for a specific multicast group to a single downstream channel in the Load Balancing Group. More generally, a load balancing algorithm will perform more effectively if it takes into account both the unicast and multicast traffic load for each CM when making decisions on where and when to move CMs.

With CMs performing Multicast DSID Forwarding (clause 9.2), the CMTS is aware of each IP multicast session joined by CPEs behind a CM. In this case, the CMTS can maintain proper IP multicast replication when autonomously moving the received downstream channels of a CM. This is not always the case for CMs not performing Multicast DSID Forwarding, where the CMTS may be unaware of which CMs have multicast group members and which do not.

CMs not performing Multicast DSID Forwarding track IGMP messages in order to control multicast group forwarding state. The IGMP protocol requires hosts to suppress IGMP messages that are not necessary for the router to maintain multicast group membership state. The [14] and [13] specifications and clause G.4.3.2, extend these IGMP requirements to the DOCSIS access network by requiring CMs to suppress messages that are deemed to be superfluous for the CMTS. As a result, the CMTS is not guaranteed to be aware of multicast group membership on a per-CM basis for CMs not performing Multicast DSID Forwarding. For an active multicast group, there could be any number of CMs that have group members and that are actively forwarding multicast traffic, but that have not sent a Membership Report to the CMTS. This lack of CMTS awareness can create a situation in which load balancing and multicast conflict.

If a CM with active multicast sessions is moved from its current downstream to a new downstream that is not carrying the multicast traffic, the session will be interrupted until the CM or CPE sends a Membership Report. In order to reduce the interruption of multicast service, CMs that implement active IGMP mode (see clause G.4.3.2) are recommended to send a Membership Report for all active multicast groups upon completion of a DCC or DBC operation that involves a downstream channel change.

The multicast issues are alleviated to some degree when BPI+ is enabled and are alleviated further when multicast traffic is encrypted using dynamic security associations (see [15]).

When BPI+ is enabled, a CM will, upon receiving an IGMP "join" message on its CPE interface, send an SA Map Request message to the CMTS. Since this message is only sent at the moment multicast group membership begins, it does not provide any indication of ongoing membership. Because multicast group membership can be transient, the past receipt of an SA Map Request for a particular multicast group, although necessary, is not a sufficient condition to alert the CMTS that the CM currently has members for that multicast group. The absence of an SA Map Request is sufficient evidence that the CM does not have members for the multicast group.

If the multicast traffic for a particular multicast group is encrypted using a dynamic security association, the CMTS can monitor the reception of TEK Key Requests and gain knowledge of multicast group membership. Since it is optional functionality for a CM to stop the TEK state machine (and discontinue sending Key Requests) when there are no longer members for multicast groups mapped to a particular security association, the continued receipt of Key Requests by the CMTS does not necessarily indicate continued multicast group membership. The lack of continuing Key Requests, however, does indicate lack of members.

11.6.8 Externally-Directed Load Balancing

The CMTS **MUST** support a means (via CMTS management objects) for an operator or external entity to direct the CMTS to initiate a DCC or DBC transaction with a CM. Due to the potential conflict between this functionality and the algorithms of the CMTS's own Autonomous Load Balancing functionality, the CMTS **MAY** reject such directions when Autonomous Load Balancing is enabled.

12 Supporting Future New Cable Modem Capabilities

12.1 Downloading Cable Modem Operating Software

A CMTS **SHOULD** be capable of being remotely reprogrammed in the field via a software download over 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 CM **MUST** implement a TFTP client compliant with [31] for software file downloads. The CM **MAY** implement an HTTP client compliant with [35] or [44] for software file downloads. The transfer is SNMP-initiated, as described in [10] or configuration file-initiated, as described here.

If the file specified in the configuration file SW Upgrade File Name TLV does not match the current software image of the CM, the CM **MUST** request the specified file via TFTP from the software server. The CM selects the software server as follows:

- If the CM downloads via IPv4 a configuration file which includes the Software Upgrade IPv4 TFTP Server TLV, the CM **MUST** use the server specified by this TLV. If the CM downloads via IPv4 a configuration file which does not include the Software Upgrade IPv4 TFTP Server TLV, the CM **MUST** use the IPv4 TFTP server from which it downloaded the configuration file. The CM **MUST** ignore the Software Upgrade IPv6 TFTP Server TLV when it downloads a configuration file using IPv4.
- If the CM downloads via IPv6 a configuration file which includes the Software Upgrade IPv6 TFTP Server TLV, the CM **MUST** use the server specified by this TLV. If the CM downloads via IPv6 a configuration file which does not include the Software Upgrade IPv6 TFTP Server TLV, the CM **MUST** use the IPv6 TFTP server from which it downloaded the configuration file. The CM **MUST** ignore the Software Upgrade IPv4 TFTP Server TLV when it downloads a configuration file using IPv6.

The CM performs the download after it registers and, if BPI is enabled, after it initializes baseline privacy. When performing a configuration-file-initiated software download, the CM **MAY** defer bridging between the RF and CPE ports until the download is complete. 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. The CM **MAY** report the failure 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.

If the CM receives an ICMP Destination Unreachable message or ICMP port unreachable message for the TFTP server at any time during the firmware download process, the CM **MUST** terminate the firmware download on the TFTP server whose address is included in the ICMP Destination Unreachable message without performing the TFTP Read Request Retries or the TFTP Download Retries (annex B).

12.2 Future Capabilities

If the CM indicates support for one or more CM capabilities defined in a higher-numbered version of DOCSIS, it **MUST** implement them in a manner that complies with the specification in which the feature is defined.

Annex A (normative): Well_known_Addresses

A.1 Addresses

A.1.1 General MAC Addresses

MAC addresses described here are defined using the Ethernet/ISO8802-3 [18] convention as bit-little-endian.

The CMTS MUST use the "All CMs Multicast MAC Address" to address the set of all CMs; for example, when transmitting Allocation Map PDUs. The CM MUST accept all traffic received with the "All CMs Multicast MAC Address".

- All CMs Multicast MAC Address: 01-E0-2F-00-00-01.

The addresses in the range:

- Reserved Multicast MAC Addresses: 01-E0-2F-00-00-02 through 01-E0-2F-00-00-0F.

are reserved for future definition. Frames addressed to any of the "Reserved Multicast MAC Addresses" SHOULD NOT be forwarded by the CM. Frames addressed to any of the "Reserved Multicast MAC Addresses" SHOULD NOT be forwarded by the CMTS.

A.1.2 Well-known IPv6 Addresses

IPv6 networks communicate using several well-known addresses per [i.7] described in table A-1.

Table A-1: Well-known IPv6 Addresses

Well-known IPv6 MAC Addresses	Well-known IPv6 Addresses	Description
33-33-00-01-00-02	FF02::1:2	All DHCP relay agents and servers
33-33-00-01-00-03	FF05::1:3	All DHCP servers
33-33-FF-xx-xx-xx	FF02:0:0:0:1:FFxx:xxxx	Link-local scope solicited node multicast address
33-33-00-00-00-02	FF02::2	Link-local scope all routers multicast address
33-33-00-00-00-01	FF02::1	Link-local scope all nodes multicast address

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

The following Service IDs are used in MAPs for special purposes or to indicate that any CM can respond in the corresponding interval.

- 0x0000 is addressed to no CM. This address is typically used when changing upstream burst parameters so that CMs have time to adjust their modulators before the new upstream settings take effect. The CM MUST NOT transmit during any transmit opportunity that has been assigned to this SID. This is also the "Initialization SID" used by the CM during initial ranging.

- 0x3FFF is addressed to all CMs. It is typically used for broadcast Request intervals or broadcast Initial Maintenance intervals.

A.2.2 Well-Known Multicast Service IDs

The following Service IDs are only used for Request/Data IEs. They indicate that any CM can respond in a given interval, but that the CM must limit the size of its transmission to a particular number of mini-slots (as indicated by the particular multicast SID assigned to the interval).

0x3FF1-0x3FFE is addressed to all CMs. IDs in this range are 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

The following Service IDs (0x3Exx) are reserved for Request IEs (refer to clause 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

The CMTS MUST carry all DOCSIS data in MPEG-2 packets with the header PID field set to 0x1FFE.

Annex B (normative): Parameters and Constants

Table B-1: Parameters and Constants

System	Name	Parameter Description	Minimum Value	Default Value	Maximum Value
CMTS	Sync Interval	Nominal time between transmission of SYNC messages (refer to clause 6.4.2).			200 ms
CMTS	UCD Interval	Time between transmission of UCD messages (refer to clause 6.4.3).			2 s
CMTS	Max MAP Pending	The number of mini-slots that a CMTS is allowed to map into the future (refer to clause 7.2.1.6).			4 096 mini-slot times
CMTS	Ranging Interval	Time between transmission of broadcast Initial Maintenance opportunities (refer to clause 7.1.3).			2 s
CM	Lost Sync Interval	Time since last received SYNC message before synchronization is considered lost.			600 ms
CM	Contention Ranging Retries	Number of Retries on Ranging Requests sent in broadcast maintenance opportunities.	16		
CM, CMTS	Invited Ranging Retries	Number of Retries on Ranging Requests sent in unicast maintenance opportunities (refer to clause 10.2.3.7).	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 (refer to clause 7.2.1.6 and "Relative Processing Delays" [12]).	(600 + M/5,12) μ s for operation in MTC mode (200 + M/5,12) μ s for operation not in MTC mode		
CMTS	CM Ranging Response processing time	Minimum time allowed for a CM following receipt of a ranging response before it is expected to transmit a ranging request in a unicast opportunity.	1 ms		
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 s		
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 ms	200 ms	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. The T4 multiplier may be set in the RNG-RSP message.	30 s (T4 Multiplier of 1)	30 s	300 s (T4 Multiplier of 10)
CMTS	T5	Wait for Upstream Channel Change response.			2 s
CM CMTS	T6	Wait for REG-RSP, REG-RSP-MP, or REG-ACK.		3 s	

System	Name	Parameter Description	Minimum Value	Default Value	Maximum Value
CM CMTS	Mini-slot size for 1.x channels.	Size of mini-slot for upstream transmission. For channels that support DOCSIS 1.x CMs.	32 modulation intervals		
CM CMTS	Mini-slot size for DOCSIS 2.0 Only Channels.	Size of mini-slot for upstream transmission. For channels that do not support DOCSIS 1.x CMs.	16 symbols		
CM CMTS	Timebase Tick	System timing unit.	6,25 μ s		
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 s
CM CMTS	T8	Wait for DSA/DSC Acknowledge timeout.			300 ms
CM	TFTP Backoff Start	Initial value for TFTP backoff.	1 s		
CM	TFTP Backoff End	Last value for TFTP backoff.	16 s		
CM	TFTP Request Retries	Number of retries on TFTP request.	4		
CM	TFTP Download Retries	Number of retries on entire TFTP downloads.	3		
CM	TFTP Wait	The wait between TFTP retry sequences.	3 min		
CMTS	T9	Registration Timeout, the time allowed between the CMTS sending a RNG-RSP (success) to a CM and receiving a REG-REQ or REG-REQ-MP from that same CM.	15 min	15 min	
CM CMTS	T10	Wait for Transaction End timeout.	3 s		
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 s
CM	T14	Minimum time after a DSx reject-temp-DCC and the next retry of DSx command.	2 s		
CMTS	T15	Maximum holding time for QoS resources for DCC on the new channel.	2 s		35 s
CM	T16	Maximum length of time CM remains in test mode after receiving TST-REQ message.			30 min
CM	T17	Maximum Time that CM MUST inhibit transmissions on a channel in response to its Ranging Class ID matching a bit value in the Ranging Hold-Off Priority Field.	300 s		
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 s
CM	DCI-REQ retry	Number of retries of DCI-REQ before rebooting.			16
CM	DCI Backoff start	Initial value for DCI backoff.	1 s		
CM	DCI Backoff end	Last value for DCI backoff.	16 s		
CMTS	CM UCD processing time	Time between the transmission of the last bit of a UCD with a new Change Count and the transmission time of the first bit of the first MAP using the new UCD (see clause 11.1).	1,5 ms * The number of upstream channels modified simultaneously		

System	Name	Parameter Description	Minimum Value	Default Value	Maximum Value
CMTS	DBC-REQ Retries	Maximum number of times the CMTS will retransmit a DBC-REQ while awaiting the DBC-RSP from the CM.	6		
CM	DBC-REQ Timeout	The amount of time that the CM waits to receive all fragments of the DBC-REQ message.	1 s		
CM	DBC-RSP Retries	Maximum number of times the CM will retransmit a DBC-RSP while awaiting the DBC-ACK from the CMTS.	6		
CM	DBC-ACK timeout	The amount of time that the CM waits for DBC-ACK after sending DBC-RSP.	300 ms		
CM	DBC DS Acquisition timeout	The amount of time that the CM is to continue trying to acquire downstream channels added to the RCS in a DBC-REQ message.	1 s		
CMTS	Sequence Hold timeout	The time that the CMTS waits before changing the Sequence Change Count for a resequencing DSID.	1 s		
CM	DSID filter count	The total number of DSID filters clause 6.2.5.6.	32		
CM	DSID resequencing context count	The number of DSIDs for re-sequencing.	16		
CMTS	CMTS Skew Limit	Maximum interval between CMTS start of transmission of out-of-order sequenced packets on different Downstream Channels, measured at the set of CMTS [4] and [2] interfaces.		3 ms	5 ms
CM	DSID Resequencing Wait Time	Per-DSID value for the minimum interval a CM delays forwarding of a higher-numbered sequenced packet while awaiting the arrival of a lower-numbered sequenced packet.		18 ms	18 ms
CMTS	MDD Interval	Time between MDD messages on a given channel.			2 s
CM	Lost MDD timeout	Time to wait for a MDD before declaring MDD loss.	3 * Maximum MDD Interval		
CM	Initializing channel timeout CM	This field defines the maximum total time that the CM can spend performing initial ranging on the upstream channels described in the TCC of a REG-RSP, REG-RSP-MP or a DBC-REQ.		60 s	
CMTS	Initializing channel timeout CMTS	This field defines the maximum total time that the CMTS waits for a REG-ACK after sending a REG-RSP-MP or waiting for a DBC-RSP after sending a DBC-REQ before retransmitting the REG-RSP-MP or DBC-REQ.		Initializing Channel Timeout CM + 3 s	
CM	T18	This timer is started when the CM receives the first Registration Response and controls the amount of time the CM waits to possibly receive a duplicate REG-RSP-MP if the REG-ACK is lost.		Initializing Channel Timeout CM + 6 s	

Annex C (normative): Common TLV Encodings

Table C-1 provides a summary of the top-level TLV encodings and the messages in which they can appear. Cfg File indicates that a particular TLV is intended to appear in the CM configuration file. REG indicates that a particular TLV can appear in at least one of the following messages: REG-REQ, REG-REQ-MP, REG-RSP, REG-RSP-MP or REG-ACK. DSx indicates that a particular TLV can appear in at least one of the following messages: DSA-REQ, DSA-RSP, DSA-ACK, DSC-REQ, DSC-RSP, DSC-ACK. DBC indicates that a particular TLV can appear in at least one of the following messages: DBC-REQ, DBC-RSP, DBC-ACK. This table is informative, detailed requirements for the placement of these TLVs in different messages are provided in the referenced clauses.

Table C-1: Summary of Top-Level TLV Encodings

Type	Description	Length	Cfg File	REG	DSx	DBC	Clause
0	Pad	-	x				C.1.2.2
1	Downstream Frequency	4	x	x			C.1.1.1
2	Upstream Channel ID	1	x	x			C.1.1.2
3	Network Access Control Object	1	x	x			C.1.1.3
4	DOCSIS 1.0 Class of Service	n	x	x			C.1.1.4
5	Modem Capabilities	n		x			C.1.3.1
6	CM Message Integrity Check (MIC)	16	x	x			C.1.1.5
7	CMTS Message Integrity Check (MIC)	16	x	x			C.1.1.6
8	Vendor ID Encoding	3		x			C.1.3.2
9	SW Upgrade Filename	n	x				C.1.2.3
10	SNMP Write Access Control	n	x				C.1.2.4
11	SNMP MIB Object	n	x				C.1.2.5
12	Modem IP Address	4		x			C.1.3.3
13	Service(s) Not Available Response	3		x			C.1.3.4
14	CPE Ethernet MAC Address	6	x				C.1.2.6
15	Telephone Settings Option (deprecated)						
16	Void						
17	Baseline Privacy	n	x	x			C.3.1
18	Max Number of CPEs	1	x	x			C.1.1.7
19	TFTP Server Timestamp	4	x	x			C.1.1.8
20	TFTP Server Provisioned Modem IPv4 Address	4	x	x			C.1.1.9
21	SW Upgrade IPv4 TFTP Server	4	x				C.1.2.7
22	Upstream Packet Classification	n	x	x	x		C.1.1.11/C.2.1.1
23	Downstream Packet Classification	n	x	x	x		C.1.1.12/C.2.1.3
24	Upstream Service Flow	n	x	x	x		C.1.1.13/C.2.2.1
25	Downstream Service Flow	n	x	x	x	x	C.1.1.14/C.2.2.2
26	Payload Header Suppression	n	x	x	x	x	C.1.1.15/C.2.2.8
27	HMAC-Digest	20			x	x	C.1.4.1
28	Maximum Number of Classifiers	2	x	x			C.1.1.16
29	Privacy Enable	1	x	x			C.1.1.17
30	Authorization Block	n			x		C.1.4.2
31	Key Sequence Number	1			x	x	C.1.4.3
32	Manufacturer Code Verification Certificate	n	x				C.1.2.10
33	Co-Signer Code Verification Certificate	n	x				C.1.2.11
34	SNMPv3 Kickstart Value	n	x				C.1.2.9
35	Subscriber Mgmt Control	3	x	x			C.1.1.19.1
36	Subscriber Mgmt CPE IPv4 List	n	x	x			C.1.1.19.2
37	Subscriber Mgmt Filter Groups	8	x	x			C.1.1.19.4
38	SNMPv3 Notification Receiver	n	x				C.1.2.12
39	Enable 2.0 Mode	1	x	x			C.1.1.20
40	Enable Test Modes	1	x	x			C.1.1.21
41	Downstream Channel List	n	x	x			C.1.1.22
42	Static Multicast MAC Address	6	x				C.1.1.23
43	DOCSIS Extension Field	n	x	x			C.1.1.18
44	Vendor Specific Capabilities	n		x			C.1.3.5

Type	Description	Length	Cfg File	REG	DSx	DBC	Clause
45	Downstream Unencrypted Traffic (DUT) Filtering	n	x	x			C.1.1.24
46	Transmit Channel Configuration (TCC)	n		x		x	C.1.5.1
47	Service Flow SID Cluster Assignment	n		x	x	x	C.1.5.2
48	Receive Channel Profile	n		x			C.1.5.3.1
49	Receive Channel Configuration	n		x		x	C.1.5.3.1
50	DSID Encodings	n		x		x	C.1.5.3.8
51	Security Association Encoding	n		x		x	C.1.5.5
52	Initializing Channel Timeout	2		x		x	C.1.5.6
53	SNMPv1v2c Coexistence	n	x				C.1.2.13
54	SNMPv3 Access View	n	x				C.1.2.14
55	SNMP CPE Access Control	1	x				C.1.2.15
56	Channel Assignment	n	x	x			C.1.1.25
57	CM Initialization Reason	1		x			C.1.3.6
58	SW Upgrade IPv6 TFTP Server	16	x				C.1.2.8
59	TFTP Server Provisioned Modem IPv6 Address	16	x	x			C.1.1.10
60	Upstream Drop Packet Classification	n	x	x	x		C.2.1.2
61	Subscriber Mgmt CPE IPv6 Prefix List	n	x	x			C.1.1.19.3
62	Upstream Drop Classifier Group ID	n	x	x			C.1.1.26
63	Subscriber Mgmt Control Max CPE IPv6 Addresses	n	x	x			C.1.1.19.5
64	CMTS Static Multicast Session Encoding	n	x				C.1.1.27
65	L2VPN MAC Aging Encoding	n	x				[8]
66	Management Event Control Encoding	n	x				C.1.2.16
67	Subscriber Mgmt CPE IPv6 List	n	x	x			C.1.1.19.6
255	End-of-Data	-	x				C.1.2.1

C.1 Encodings for Configuration and MAC-Layer Messaging

The following type/length/value encodings **MUST** be used by CMs and CMTSs in both the configuration file (see annex 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

The TLVs in the following clauses are intended to be forwarded by the CM to the CMTS in the Registration Request message. Some of these TLVs require inclusion in the E-MIC Bitmap in order to be utilized by the CMTS.

C.1.1.1 Downstream Frequency Configuration Setting

The frequency of the Primary Downstream Channel to be used by the CM for initialization unless a Downstream Channel List is present in the configuration file. It is an override for the CM's Primary Downstream 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 62 500 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, CPEs 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 continue to accept and generate traffic from the CM itself and not forward traffic from an attached CPE to the RF MAC Network. The value of this field does not affect CMTS service flow operation and does not affect CMTS data forwarding operation.

Type	Length	Value
3	1	1 or 0

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 "trace-route".
- 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.
- HTTP (if supported): allow the modem to download new a software image.

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 is treated by the CMTS as described in clause 6.4.8.3.2.

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. The CM **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 1: 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:

- 1) Discarding over-limit requests.
- 2) Deferring (through zero-length grants) the grant until it is conforming to the allowed limit.
- 3) Discarding over-limit data packets.
- 4) Reporting to a policy monitor (for example, using the alarm mechanism) that is capable of incapacitating errant CMs.

NOTE 2: 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 to 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 [15].

Type	Length	Enable / Disable
4.7 (= CoS_BP_ENABLE)	1	1 or 0

Table C-2: 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/s
		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/s
		6	2	1 518	max. Tx burst of 1 518 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/s
		3	4	300 000	max. return rate of 300 Mb/s
		4	1	3	return path priority of 3
		5	4	32 000	min guaranteed 32 kb/s
		6	2	1 518	max. Tx burst of 1 518 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. The length of this value field is a function of the Extended CMTS MIC HMAC type (an MD5 HMAC requires 16 bytes; other HMAC types may produce longer or shorter digests). HMAC types which produce a digest of fewer than 16 bytes MUST be padded with zeros to 16 bytes.

Type	Length	Value
7	$n \geq 16$	d1, d2,... d16,... dn

C.1.1.7 Maximum Number of CPEs

The maximum number of CPEs which can be granted access through a CM during a CM epoch. The CM epoch is the time between startup and hard reset of the modem. The maximum number of CPEs MUST be enforced by the CM.

NOTE 1: 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 clause 9.1.2.1). 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 by the CM as the default value of 1.

NOTE 2: 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 [27].

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 IPv4 Address

The IPv4 Address of the modem requesting the configuration file.

Type	Length	Value
20	4	IPv4 Address

NOTE: The purpose of this parameter is to prevent IP spoofing during registration.

C.1.1.10 TFTP Server Provisioned Modem IPv6 Address

The IPv6 Address of the modem requesting the configuration file.

Type	Length	Value
59	16	IPv6 Address

NOTE: The purpose of this parameter is to prevent IP spoofing during registration.

C.1.1.11 Upstream Packet Classification Configuration Setting

This field defines the parameters associated with one entry in an upstream traffic classification list. Refer to clause C.2.1.1.

Type	Length	Value
22	N	

C.1.1.12 Downstream Packet Classification Configuration Setting

This field defines the parameters associated with one Classifier in a downstream traffic classification list. Refer to clause C.2.1.3.

Type	Length	Value
23	N	

C.1.1.13 Upstream Service Flow Encodings

This field defines the parameters associated with upstream scheduling for one Service Flow. Refer to clause C.2.2.1.

Type	Length	Value
24	N	

C.1.1.14 Downstream Service Flow Encodings

This field defines the parameters associated with downstream scheduling for one Service Flow. Refer to clause C.2.2.2.

Type	Length	Value
25	N	

C.1.1.15 Payload Header Suppression

This field defines the parameters associated with Payload Header Suppression.

Type	Length	Value
26	N	

C.1.1.16 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. Upstream Drop Classifiers are not 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 used by the CM and CMTS MUST be 0 - no limit.

C.1.1.17 Privacy Enable

This configuration setting enables/disables Baseline Privacy [15] on the Primary Service Flow and all other Service Flows for this CM. If a DOCSIS 2.0 or 3.0 CM receives this setting in a configuration file, the CM is required to forward this setting as part of the Registration Request (REG-REQ or REG-REQ-MP) as specified in clause 6.4.7, 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 used by the CM and CMTS MUST be 1 - privacy enabled.

C.1.1.18 DOCSIS Extension Field

The DOCSIS Extension Field is used to extend the capabilities of the DOCSIS specification, through the use of new and/or vendor-specific features.

The DOCSIS Extension Field must be encoded using TLV 43 and must include the Vendor ID field (refer to clause C.1.3.2) to indicate whether the DOCSIS Extension Field applies to all devices or only to devices from a specific vendor. The Vendor ID must be the first TLV embedded inside the DOCSIS Extension Field. If the first TLV inside the DOCSIS Extension Field is not a Vendor ID, then the TLV MUST be discarded by the CMTS. In this context, the Vendor ID of 0xFFFFF is reserved to signal that this DOCSIS Extension Field contains general extension information (see clause C.1.1.18.1); otherwise, the DOCSIS Extension Field contains vendor-specific information (see clause C.1.1.18.2).

This configuration setting 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. The same Vendor ID may appear multiple times. However, there must not be more than one Vendor ID TLV inside a single TLV 43.

The CM MUST ignore any DOCSIS Extension Field that it cannot interpret, but still include the TLV in the REG-REQ or REG-REQ-MP message. The CM MUST NOT initiate the DOCSIS Extension Field TLVs.

Type	Length	Value
43	N	

C.1.1.18.1 General Extension Information

When using the DOCSIS Extension Field (TLV 43) to encode general extension information, the Vendor ID of 0xFFFFF must be used as the first sub-TLV inside TLV 43.

Type	Length	Value
43	N	8, 3, 0xFFFFF, followed by general extension information

The following sub-TLVs are defined only as part of the General Extension Information. The type values may be re-defined for any purpose as part of a Vendor-Specific Information encoding.

C.1.1.18.1.1 CM Load Balancing Policy ID

The CMTS load balancing algorithm uses this config file setting as the CM load balancing policy id. If present, this value overrides the default group policy assigned by the CMTS (see clause 11.6). This configuration setting should only appear once in a configuration file. This configuration setting must only be used in configuration files, REG-REQ and REG-REQ-MP messages and must not be nested inside a Packet Classification Configuration Setting, a Service Flow Configuration Setting or a Service Flow Response.

Type	Length	Value
43.1	4	Policy id

C.1.1.18.1.2 CM Load Balancing Priority

This config file setting is the CM load balancing priority to be used by the CMTS load balancing algorithm. If present, this value overrides the default priority assigned by the CMTS (see clause 11.6). This configuration setting should only appear once in a configuration file. This configuration setting must only be used in configuration files, REG-REQ and REG-REQ-MP messages and must not be nested inside a Packet Classification Configuration Setting, a Service Flow Configuration Setting or a Service Flow Response.

Type	Length	Value
43.2	4	Priority

C.1.1.18.1.3 CM Load Balancing Group ID

This config file setting is the Restricted Load Balancing Group ID defined at the CMTS. If present, this value overrides the general load balancing group. If no Restricted Load Balancing Group is defined that matches this group id, the value is ignored by the CMTS (see clause 11.6). This configuration setting should only appear once in a configuration file. This configuration setting must only be used in configuration files, REG-REQ and REG-REQ-MP messages and must not be nested inside a Packet Classification Configuration Setting, a Service Flow Configuration Setting or a Service Flow Response.

Type	Length	Value
43.3	4	Group id

C.1.1.18.1.4 CM Ranging Class ID Extension

This config file setting is the CM Ranging Class ID Extension to be defined by the cable operator. These bits will be prepended to the CM's default Ranging Class ID as the most significant bits of the 32 bit Ranging Class ID value. These bits will be sent in the REG-REQ or REG-REQ-MP as part of the CM's Ranging Class ID in the modem capabilities field. If the TLV is not included in the configuration file, the CM will use zero for this value. These bits allow the user to define special device classes that could be used to give those devices or service types, preferential treatment with respect to ranging after a massive outage. After successful registration, the CM MUST store the entire 32 bit value in non-volatile memory and use it for ranging decisions after a reboot or a re-init MAC event.

Type	Length	Value
43.4	2	Extended ID

C.1.1.18.1.5 L2VPN Encoding

The L2VPN Encoding parameter is a multi-part encoding that configures how the CMTS performs Layer 2 Virtual Private Network bridging for CPE packets. The subtypes of the L2VPN encoding are specified in [8]. The CMTS MAY support the DOCSIS Layer 2 Virtual Private Network feature as defined in [8]. If the L2VPN feature is not supported, the CMTS MUST ignore the information in the L2VPN configuration setting.

Type	Length	Value
43.5	n	L2VPN Encoding subtype/length/value tuples

C.1.1.18.1.6 Extended CMTS MIC Configuration Setting

The Extended CMTS MIC Configuration Setting parameter is a multi-part encoding that configures how the CMTS performs message integrity checking. This is used to detect unauthorized modification or corruption of the CM configuration file, using techniques which are not possible using the pre-3.0 DOCSIS CMTS MIC, in particular, using more advanced hashing techniques or requiring different TLVs to be included in the HMAC calculation. This TLV cannot be contained within an instance of TLV type 43 which contains other subtypes (excluding subtype 8).

Type	Length	Value
43.6	n	Extended CMTS MIC Parameter Encoding subtype/length/value tuples

C.1.1.18.1.6.1 Extended CMTS MIC HMAC type

The Extended CMTS MIC HMAC type parameter is a single byte encoding that identifies the type of hashing algorithm used in the CMTS MIC hash TLV. This subtype is always included within an Extended CMTS MIC Configuration Setting TLV; the instance of the CMTS MIC Hash within the configuration file will use the HMAC technique described by the value of this TLV.

The CMTS SHOULD support a configuration that can require all REG-REQ or REG-REQ-MP messages to contain an Extended CMTS MIC Encoding with a particular CMTS MIC algorithm.

Type	Length	Value
43.6.1	1	Enumeration 1 - MD5 HMAC [36] 2 - MMH16-σ-n HMAC [15] 43 - vendor specific

C.1.1.18.1.6.2 Extended CMTS MIC Bitmap

The Extended CMTS MIC Bitmap is a multi byte encoding that is a bitmask representing specified TLV types in a CM configuration file, REG-REQ or REG-REQ-MP message, clause D.2. This TLV is always present and the TLVs to be included within the digest calculation are those whose top level types correspond to bits which are set in this value. For example, to require the Downstream Frequency Configuration Setting to be included in the digest calculation, set bit 1 in the value of this TLV. This TLV uses the BITS Encoding convention where bit positions are numbered starting with bit #0 as the most significant bit.

Type	Length	Value
43.6.2	n	BITS Encoding - Each bit position in this string represents a top-level TLV Bit position 0 is reserved and is always set to a value of 0

C.1.1.18.1.6.3 Explicit Extended CMTS MIC Digest Subtype

This subtype explicitly provides the calculated extended MIC digest value over all TLVs reported in REG-REQ or REG-REQ-MP for which bits are set in the Extended CMTS MIC Bitmap. If the Extended CMTS MIC Bitmap indicates TLV 43 is to be included in the calculation of the Extended CMTS MIC digest, this subtype (with the value 0) is to be included in that calculation, see clauses D.1.3 and D.2.1. A valid Explicit Extended CMTS MIC Digest does NOT contain the CM MIC value.

When this subtype is present, the CMTS MIC Configuration Setting in TLV7 is calculated using the set of TLVs as specified for DOCSIS 2.0, in clause D.2.1.

If this subtype is omitted from an Extended CMTS MIC Encoding, the extended CMTS MIC is implicitly provided in the CMTS MIC Configuration Setting of TLV 7.

When the Explicit Extended CMTS MIC Digest Subtype is present, if the CMTS fails the Extended CMTS MIC Digest verification but passes the pre-3.0 DOCSIS CMTS MIC digest verification of TLV7, then the CMTS MUST NOT consider the CM to have failed authentication. Instead, the CMTS MUST silently ignore all TLVs in REG-REQ or REG-REQ-MP which were marked as protected by the Extended CMTS MIC Bitmap but are not included in the set of TLVs protected by the pre-3.0 DOCSIS CMTS MIC (TLV7) calculation.

Type	Length	Value
43.6.3	n	Calculated MIC digest using the CMTS MIC HMAC Type algorithm

C.1.1.18.1.7 Source Address Verification (SAV) Authorization Encoding

This parameter configures a static range of IP addresses authorized for the Source Address Verification (SAV) enforced by the CMTS for upstream traffic from the CM (see [15]). It is intended to be configured for CMs connecting to CPEs with statically configured CPE Host IP addresses or for CMs connecting to a customer premise IP router that reaches a statically assigned IP subnet.

This parameter is intended for the CMTS only and is ignored by the CM. The parameter is encoded as a subtype of the DOCSIS Extension Information TLV43 encoding in order for it to be included by CMs supporting any DOCSIS version.

An IP address "prefix" is a combination of an IP address (the "prefix address") and a bit count (the "prefix length"). An IP address is said to be "within" a prefix when it matches the prefix length number of most significant bits in the prefix address. A prefix length of zero means that all IP addresses are within the prefix.

The SAV Authorization Encoding defines either or both of:

- a "SAV Group Name" that indirectly identifies an "SAV Group", which is a configured list of prefixes in the CMTS; or
- a list of "Static SAV Prefix Rules", each of which directly defines a single prefix.

The CMTS considers an upstream source IP address within any of the above mentioned prefixes to be authorized for purposes of Source Address Verification.

A valid configuration file, REG-REQ or REG-REQ-MP message contains at most one instance of the SAV Authorization Encoding. Other restrictions on the subtypes of a valid SAV Authorization Encoding are described below. CM and CMTS operation with an invalid SAV Authorization Encoding is not specified.

Type	Length	Value
43.7	N	Subtype encodings

C.1.1.18.1.7.1 SAV Group Name Subtype

This subtype contains an ASCII string that identifies an SAV Group Name configured in the CMTS.

Type	Length	Value
43.7.1	1..15	Name of an SAV Group configured in the CMTS

A valid SAV Authorization Encoding contains zero or one instances of this subtype.

A CMTS MUST support registration of CMs that reference an SAV Group Name that does not exist in the CMTS. A CMTS MUST support creation, modification and deletion of configured SAV Groups while CMs remain registered that reference the SAV Group Name.

C.1.1.18.1.7.2 SAV Static Prefix Rule Subtype

This subtype identifies a single static prefix within which upstream traffic from the CM is authorized for purposes of Source Address Verification. A valid SAV Authorization Encoding contains zero, one or more instances of this subtype. A CMTS MUST support at least one SAV Static Prefix Rule for each CM.

The CMTS maintains a management object that reports for each CM the list of SAV Static Prefixes learned from that CM in its REG-REQ or REG-REQ-MP. The CMTS is expected to recognize when multiple CMs report the same list of SAV Static Prefix Rules. The CMTS assigns a "list identifier" to each unique set of SAV prefixes. The minimum number of different SAV Static Prefix lists supported by a CMTS is vendor-specific.

Type	Length	Value
43.7.2	N	SAV Static Prefix Subtype encodings

C.1.1.18.1.7.2.1 SAV Static Prefix Address Subtype

This subtype identifies an IPv4 or IPv6 address subnet authorized to contain a source IP address of upstream traffic. A valid SAV Static Prefix Rule Subtype contains exactly one instance of this subtype.

Type	Length	Value
43.7.2.1	4 (IPv4) or 16 (IPv6)	Prefix of an IP address range authorized to contain the source IP address for upstream packets

C.1.1.18.1.7.2.2 SAV Static Prefix Length Subtype

This subtype defines the number of most significant bits in an SAV Static Prefix Address. A valid SAV Static Prefix Rule Subtype contains exactly one instance of this subtype.

Type	Length	Value
43.7.2.2	1	Range 0..32 for an IPv4 SAV Static Prefix Address or 0..128 for an IPv6 SAV Static Prefix Address. Number of most significant bits of the Static SAV Prefix Address matched to an upstream source IP address. A value of 0 means that all source addresses are authorized for SAV

C.1.1.18.1.8 Cable Modem Attribute Masks

If specified, this TLV limits the set of channels to which the CMTS SHOULD assign the cable modem by requiring or forbidding certain binary attributes. This TLV is primarily intended for CMs not operating in Multiple Receive Channel mode. It is CMTS vendor-specific whether or not this TLV is used in channel assignment for CMs operating in Multiple Receive Channel mode. When Service Flow Attribute Masks are present in the CM configuration file as well, the CMTS will observe the precedence order defined in clause 10.2.6.2.1.

See clause 8.1.1 for how the Required Attribute mask, Forbidden Attribute Mask control how CMs may be assigned to particular channels.

Type	Length	Value
43.9	n	Cable Modem Attribute Mask subtype encodings

C.1.1.18.1.8.1 Cable Modem Required Downstream Attribute Mask

If specified, this sub-TLV limits the set of downstream channels to which the CMTS assigns the cable modem requiring certain binary attributes.

Type	Length	Value
43.9.1	4	32-bit mask representing the set of binary channel attributes required for the CM

C.1.1.18.1.8.2 Cable Modem Downstream Forbidden Attribute Mask

If specified, this sub-TLV limits the set of downstream channels to which the CMTS assigns the CM by forbidding certain attributes.

Type	Length	Value
43.9.2	4	32-bit mask representing the set of binary channel attributes forbidden for the CM

C.1.1.18.1.8.3 Cable Modem Upstream Required Attribute Mask

If specified, this sub-TLV limits the set of upstream channels to which the CMTS assigns the cable modem requiring certain binary attributes.

Type	Length	Value
43.9.3	4	32-bit mask representing the set of binary channel attributes required for the CM

C.1.1.18.1.8.4 Cable Modem Upstream Forbidden Attribute Mask

If specified, this sub-TLV limits the set of upstream channels to which the CMTS assigns the CM by forbidding certain attributes.

Type	Length	Value
43.9.4	4	32-bit mask representing the set of binary channel attributes forbidden for the CM

C.1.1.18.1.9 IP Multicast Join Authorization Encoding

This subtype of the DOCSIS Extension Information (TLV43) encoding identifies a set of IP Multicast Join Authorization session rules. This parameter is intended for the CMTS only and is ignored by the CM. The parameter is encoded as a subtype of the DOCSIS Extension Information TLV43 encoding in order for it to be included by CMs supporting any DOCSIS version. A CMTS uses the IP Multicast Join Authorization Encoding to authorize IP multicast session joins for all DOCSIS CM versions.

A valid CM configuration file and CM Registration Request contains zero or one instances of the IP Multicast Join Authorization Encoding. Other restrictions on the subtypes of a valid IP Multicast Join Authorization Encoding are described below. CM and CMTS operation with an invalid IP Multicast Join Authorization Encoding is not specified.

Type	Length	Value
43.10	N	IP Multicast Join Authorization Subtype encodings

C.1.1.18.1.9.1 IP Multicast Profile Name Subtype

This subtype contains an ASCII string that identifies an IP Multicast Profile Name configured in the CMTS.

Type	Length	Value
43.10.1	1..15	Name of an IP Multicast Profile configured in the CMTS

A valid IP Multicast Join Authorization Encoding contains zero, one or more instances of this subtype.

C.1.1.18.1.9.2 IP Multicast Join Authorization Static Session Rule Subtype

This subtype statically configures a single IP multicast "session rule" that controls the authorization of a range of IP multicast sessions. A session rule identifies a CMTS join authorization action of "permit" or "deny" for the combination of a range of source addresses (an "S prefix") and destination group addresses (a "G prefix") of a multicast session.

An IP address "prefix" is a combination of an IP address (the "prefix address") and a bit count (the "prefix length"). An IP address is said to be "within" a prefix when it matches the prefix length number of most significant bits in the prefix address. A prefix length of zero means that all IP addresses are within the prefix.

Type	Length	Value
43.10.2	N	IP Multicast Join Authorization Static Session Rule subtype encodings

A valid IP Multicast Join Authorization Encoding contains zero or more instances of this subtype.

C.1.1.18.1.9.2.1 RulePriority

This attribute configures the rule priority for the static session rule. A valid IP Multicast Join Authorization Static Session Rule Encoding contains exactly one instance of this subtype.

Type	Length	Value
43.10.2.1	1	0..255. Higher values indicate a higher priority. If more than one session rule matches a joined session, the session rule with the highest rule priority determines the authorization action

C.1.1.18.1.9.2.2 Authorization Action

This attribute specifies the authorization action for a session join attempt that matches the session rule. A valid IP Multicast Join Authorization Static Session Rule Encoding has exactly one instance of this subtype.

Type	Length	Value
43.10.2.2	1	0 permit 1 deny 2..255 Reserved

C.1.1.18.1.9.2.3 Source Prefix Address Subtype

This subtype identifies the prefix of a range of authorized source addresses for multicast sessions. A valid IP Multicast Join Authorization Static Session Rule Subtype contains zero or one instances of this subtype. A valid IP Multicast Join Authorization Static Session Rule Subtype either includes both a Source Prefix Address Subtype and a Source Prefix Length Subtype or omits both Source Prefix Address Subtype and Source Prefix Length subtype.

If this subtype is omitted, the session rule is considered to apply to all sources of multicast sessions.

Type	Length	Value
43.10.2.3	4 (IPv4) or 16 (IPv6)	Prefix of an IP address range for the source of IP multicast sessions

C.1.1.18.1.9.2.4 Source Prefix Length Subtype

This subtype defines the number of matched most significant bits in the Source Prefix Address Subtype in an IP Multicast Join Authorization Static Session Rule Subtype.

Type	Length	Value
43.10.2.4	1	Number of most significant bits of the Source Prefix Address matched to the source IP address of a source-specific multicast session. The value range is 0..32 for an IPv4 Source Prefix Address or 0..128 for an IPv6 Source Prefix Address. A value of 0 means that all source addresses are matched by the rule

C.1.1.18.1.9.2.5 Group Prefix Address Subtype

This subtype identifies the prefix of a range of destination IP multicast group addresses. A valid IP Multicast Join Authorization Static Session Rule Subtype contains exactly one instance of this subtype.

Type	Length	Value
43.10.2.5	4 (IPv4) or 16 (IPv6)	Prefix of an IP address range for the destination group of IP multicast sessions

C.1.1.18.1.9.2.6 Group Prefix Length Subtype

This subtype defines the number of matched most significant bits in the Group Prefix Address Subtype in an IP Multicast Join Authorization Static Session Rule Subtype. A valid IP Multicast Join Authorization Static Session Rule Subtype contains exactly one instance of this subtype.

Type	Length	Value
43.10.2.6	1	Number of most significant bits of the Group Prefix Address matched to an IP destination group address. The value range is 0..32 for an IPv4 Group Prefix Address or 0..128 for an IPv6 Group Prefix Address. A value of 0 means that all destination group addresses are matched by this rule

C.1.1.18.1.9.3 Maximum Multicast Sessions Encoding

This subtype, if included in an IP Multicast Join Authorization Encoding, configures the CMTS to limit the maximum number of multicast sessions authorized to be dynamically joined by clients reached through the CM.

Type	Length	Value
43.10.3	2 (unsigned 16 bit integer)	0 to 65 534: the maximum number of sessions permitted to be dynamically joined. A value of 0 indicates that no dynamic multicast joins are permitted. 65 535: no limit to the number of multicast sessions to be joined

C.1.1.18.1.10 Service Type Identifier

A text string identifying the type of service to which this CM is subscribed. This TLV is used by the CMTS to select the correct MAC Domain or Restricted Load Balancing Group to which the CM will be assigned. When this TLV is present in the Registration Request message, the CMTS MUST assign the CM to a MAC Domain or Restricted Load Balancing Group which offers the requested Service Type, if one is available. If no MAC Domain or Restricted Load Balancing Group is available that offers the requested Service Type, the CMTS is free to assign the CM to any available MAC Domain.

If this TLV is included in the configuration file along with the Load Balancing Group ID TLV, this TLV takes precedence. If the indicated Load Balancing Group is available to the CM and offers the requested Service Type, the CMTS MUST assign the CM to that Load Balancing Group. Otherwise, the CMTS ignores the Load Balancing Group ID TLV.

Type	Length	Value
43.11	1 to 16	Service Type Identifier

C.1.1.18.2 Vendor-Specific Information

Vendor-specific configuration information, if present, is encoded in the DOCSIS Extension Field (code 43) using the Vendor ID field (refer to clause C.1.3.2) to specify which TLV tuples apply to which vendor's products.

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.1.1.19 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.

C.1.1.19.1 Subscriber Management Control

This three byte field provides control information to the CMTS for the Subscriber Management requirements in [10]. The first two bytes represent the number of IPv4 addresses permitted behind the CM. The third byte is used for control fields.

Type	Length	Value
35	3	byte 1,2 MaxCpeIPv4 (low-order 10 bits)
		byte 3, bit 0: Active
		byte 3, bit 1: Learnable
		byte 3, bits #2-7: reserved, must be set to zero

C.1.1.19.2 Subscriber Management CPE IPv4 List

This field lists the IPv4 Addresses the CMTS uses as part of the total of the Max CPE IPv4 addresses in the Subscriber Management requirements in [10].

Type	Length	Value
36	N (multiple of 4)	lpa1, lpa2, lpa3, lpa4

C.1.1.19.3 Subscriber Management CPE IPv6 Prefix List

This field lists the provisioned CPE IPv6 Prefixes the CMTS uses as part of the total of the Max CPE IPv6 prefixes in the Subscriber Management requirements in [10].

Type	Length	Value
61	N (multiple of 17)	IP Prefix 1/length, IP Prefix 2/length, etc. Out of each 17 bytes, the first 16 define the IPv6 prefix and the 17th defines the length

C.1.1.19.4 Subscriber Management Filter Groups

The Subscriber Management MIB allows an upstream and downstream filter group to be assigned to a CM and its associated CPE and Service/Application Functional Entities (SAFEs). These filter groups are encoded in the configuration file in a single TLV as follows.

Type	Length	Value
37	8,12,16 or 20	Bytes 1, 2: docsSubMgtSubFilterDownstream group
		Bytes 3, 4: docsSubMgtSubFilterUpstream group
		Bytes 5, 6: docsSubMgtCmFilterDownstream group
		Bytes 7, 8: docsSubMgtCmFilterUpstream group
		Bytes 9, 10: docsSubMgtPsFilterDownstream group
		Bytes 11, 12: docsSubMgtPsFilterUpstream group
		Bytes 13, 14: docsSubMgtMtaFilterDownstream group
		Bytes 15, 16: docsSubMgtMtaFilterUpstream group
		Bytes 17, 18: docsSubMgtStbFilterDownstream group
		Bytes 19, 20: docsSubMgtStbFilterUpstream group

The elements: docsSubMgtSubFilterDownstream, docsSubMgtSubFilterUpstream, docsSubMgtCmFilterDownstream and docsSubMgtCmFilterUpstream, are considered mandatory elements. If the length is 16, the CMTS MUST use the docsSubMgtSubFilterDownstream and docsSubMgtFilterUpstream groups for filtering eSTB traffic. If the length is 12, the CMTS MUST use the docsSubMgtSubFilterDownstream and docsSubMgtFilterUpstream groups for filtering eSTB and eMTA traffic. If the length is 8, the CMTS MUST use the docsSubMgtSubFilterDownstream and docsSubMgtFilterUpstream groups for filtering eSTB, eMTA, ePS and eRouter traffic. If the length is greater than 20, the additional bytes MUST be ignored by the CMTS.

C.1.1.19.5 Subscriber Management Control Max CPE IPv6 Addresses

This field configures the maximum number of IPv6 addresses the CMTS allows forwarding traffic for the CM. This is the corresponding IPv6 version of the "Max Cpe IPv4" encoding of the Subscriber Management Control encoding (TLV 35).

Type	Length	Value
63	2	Low-order 10 bits

C.1.1.19.6 Subscriber Management CPE IPv6 List

This field lists the Ipv6 Addresses the CMTS uses as part of the total of the Max CPE Ipv6 addresses in the Subscriber Management requirements in [10].

Type	Length	Value
67	N (multiple of 16)	Ipv6 Address1, Ipv6 Address2,...Ipv6AddressN

C.1.1.20 Enable 2.0 Mode

This configuration setting enables/disables DOCSIS 2.0 mode for a CM registering:

- 1) with a DOCSIS 2.0 CMTS; or
- 2) CM registering with a DOCSIS 3.0 CMTS and not operating in Multiple Transmit Channel Mode.

When a CM is commanded to operate in Multiple Transmit Channel Mode according to the REG-RSP, this configuration setting does not have relevance. When a CM is not in Multiple Transmit Channel Mode, this configuration setting has relevance in that a CM has 2.0 mode enabled or not and if 2.0 mode is enabled the CM is actually operating in 2.0 mode if the upstream channel is of type 2,3 or 4.

The default value of this parameter used by the CM MUST be 1 - 2.0 Mode Enabled.

Type	Length	Value
39	1	0 - Disable
		1 - Enable

C.1.1.21 Enable Test Modes

This configuration setting enables/disables certain test modes for a CM which supports test modes. The definition of the test modes is beyond the scope of the present document.

If this TLV is not present, the default value used by the CM MUST be 0 - Test modes disabled.

Type	Length	Value
40	1	0 - Disable
		1 - Enable

C.1.1.22 Downstream Channel List

A list of receive frequencies to which the CM is allowed to tune during scanning operations. When the Downstream Channel List is provided in a configuration file, the CM MUST NOT attempt to establish communications using a downstream channel that is absent from this list unless specifically directed to do so by the CMTS. For example, the CMTS may direct the CM to use downstream channel(s) not listed in the Downstream Channel List via Registration Response, DBC Request and/or DCC Request message. When both the Downstream Channel List and the Downstream Frequency Configuration Setting (clause C.1.1.1) are included in the configuration file, the CM MUST ignore the Downstream Frequency Configuration Setting. This list can override the last operational channel stored in NVRAM as defined in clause 10.2.1. The CM MUST retain and employ this list of channels whenever the CM performs a re-initialize MAC or continue scanning operation. The CM MUST replace or remove the list by subsequent configuration file downloads. Upon power cycle, the CM MUST NOT enforce a previously learned downstream channel list. However, the CM MAY remember this list as an aid to downstream channel acquisition.

Type	Length	Value
41	N	List of Allowed Rx Frequencies

The list of allowed downstream frequencies is composed of an ordered series of sub-TLVs (Single Downstream Channel, Downstream Frequency Range and Default Scanning) as defined below. When scanning for a downstream channel (except after a power-cycle), the CM MUST scan through this ordered list and attempt to establish communications on the specified channel(s). The scanning is initialized as follows:

- If the CM is in an operational state and then undergoes a re-initialize MAC operation (except due to a DCC or a DBC), it MUST first scan the last operational frequency and then restart scanning at the beginning of the ordered list.
- If, while scanning this ordered list, the CM fails to become operational and is forced to re-initialize MAC, the CM MUST continue scanning from the next applicable frequency in the ordered list.
- If it reaches the Default Scanning TLV (TLV 41.3) in the configuration file, the CM begins its default scanning algorithm, completing initial ranging and DHCP and receiving a new configuration file via TFTP on the first valid frequency it sees. If the new configuration file does not contain TLV 41, the CM MUST continue with registration. If the new configuration file contains TLV 41, the CM MUST confirm that the frequency of the current Primary Downstream Channel is explicitly listed in the Downstream Channel List. If the current Primary Downstream Channel is not explicitly listed in the Downstream Channel List, the CM MUST NOT register on the current Primary Downstream Channel and MUST restart scanning according to the Downstream Channel List contained in the configuration file.

Upon reaching the end of the List, the CM MUST begin again with the first sub-TLV in the List. The CM MUST be capable of processing a Downstream Channel List that contains up to 16 sub-TLVs.

This configuration setting may appear multiple times. If this configuration setting appears multiple times, all sub-TLVs MUST be considered by the CM to be part of a single Downstream Channel List in the order in which they appear in the configuration file. In other words, the sub-TLVs from the first instance of this configuration setting would comprise the first entries in the ordered series; the second instance would comprise the next entries, etc.

C.1.1.22.1 Single Downstream Channel

Upon reaching this sub-TLV in the Downstream Channel List, the CM MUST attempt to acquire a downstream signal on the specified Frequency for a period of time specified by the Single Downstream Channel Timeout. If the channel is determined to be unsuitable for a Primary Downstream Channel by the CM, the CM MAY move on to the next sub-TLV in the Downstream Channel List without waiting for the Timeout to expire.

The CM MUST be capable of processing a Downstream Channel List that contains multiple Single Downstream Frequency TLVs.

Type	Length	Value
41.1	6 or 10	

C.1.1.22.1.1 Single Downstream Channel Timeout

Timeout is specified in seconds (unsigned). A value of 0 for Timeout means no time out, i.e. the CM attempts to acquire a signal on the specified Frequency and if unsuccessful moves immediately to the next sub-TLV in the Downstream Channel List. This is an optional parameter in a Single Downstream Channel TLV. If the Single Downstream Channel Timeout is omitted, the CM MUST use a default time out of 0.

Type	Length	Value
41.1.1	2	Timeout

C.1.1.22.1.2 Single Downstream Channel Frequency

Single Downstream Channel Frequency is a required parameter in each Single Downstream Channel TLV, the CM MUST ignore any Single Downstream Channel TLV which lacks this parameter. The DSFrequency must be a multiple of 62 500 Hz.

Type	Length	Value
41.1.2	4	DSFrequency

C.1.1.22.2 Downstream Frequency Range

Upon reaching this sub-TLV in the Downstream Channel List, the CM MUST begin scanning with DSFrequencyStart and progress in steps as indicated by DSFrequencyStepSize until reaching DSFrequencyEnd and then repeat for a period of time specified by the Downstream Frequency Range Timeout. If the value of Timeout is less than the time necessary for the CM to complete one full scan of all channels in the Downstream Frequency Range, the CM MUST complete one full scan and then move on to the next sub-TLV in the Downstream Channel List.

NOTE: DSFrequencyEnd may be less than DSFrequencyStart, which indicates scanning downward in frequency.

If a signal has been acquired on all available channels between DSFrequencyStart and DSFrequencyEnd (inclusive) and all channels have been determined to be unsuitable for a Primary Downstream Channel by the CM, the CM MAY move on to the next sub-TLV in the Downstream Channel List without waiting for the Timeout to expire.

The CM MUST be capable of processing a Downstream Channel List that contains multiple Downstream Frequency Range TLVs.

Type	Length	Value
41.2	18 or 22	

C.1.1.22.2.1 Downstream Frequency Range Timeout

Timeout is specified in seconds (unsigned). A value of 0 for Timeout means no time out, i.e. the CM attempts to acquire a signal once on each frequency within the defined range and if unsuccessful moves immediately to the next sub-TLV in the Downstream Channel List. This is an optional parameter in a Downstream Frequency Range TLV. If the Downstream Frequency Range Timeout is omitted, the CM MUST use a default for Timeout of 0.

Type	Length	Value
41.2.1	2	Timeout

C.1.1.22.2.2 Downstream Frequency Range Start

Downstream Frequency Range Start is a required parameter in each Downstream Frequency Range TLV; the CM MUST ignore any Downstream Frequency Range TLV which lacks this parameter. Downstream Frequency Range Start must be a multiple of 62 500 Hz.

Type	Length	Value
41.2.2	4	DSFrequencyStart

C.1.1.22.2.3 Downstream Frequency Range End

Downstream Frequency Range End is a required parameter in each Downstream Frequency Range TLV; the CM MUST ignore any Downstream Frequency Range TLV which lacks this parameter. Downstream Frequency Range End must be a multiple of 62 500 Hz.

Type	Length	Value
41.2.3	4	DSFrequencyEnd

C.1.1.22.2.4 Downstream Frequency Range Step Size

Downstream Frequency Range Step Size is a required parameter in each Downstream Frequency Range TLV; the CM MUST ignore any Downstream Frequency Range TLV which lacks this parameter. Downstream Frequency Range Step Size specifies the increments in Hz by which the CM MUST scan through the Downstream Frequency Range.

The CM MUST support a minimum Downstream Frequency Step Size of 6 000 000 Hz. The CM MAY support Downstream Frequency Step Sizes less than 6 000 000 Hz.

Type	Length	Value
41.2.4	4	DSFrequencyStepSize

C.1.1.22.3 Default Scanning

Upon reaching this sub-TLV in the Downstream Channel List, the CM MUST begin scanning according to its default scanning algorithm (which may be vendor dependent) and repeat for a period of time specified by Timeout. When the CM acquires a valid Primary Downstream Channel during default scanning, the CM completes initial ranging and DHCP and receives a new configuration file via TFTP. If the configuration file does not contain TLV 41, the CM continues with registration. If the configuration file contains TLV 41 and the current downstream channel is not explicitly listed in the Downstream Channel List, the CM restarts scanning according to the Downstream Channel List contained in the configuration file.

Timeout is specified in seconds (unsigned). If the value of Timeout is less than the time necessary for the CM to complete one full scan of all channels in the default scanning algorithm, the CM MUST complete one full scan and move on to the next sub-TLV in the Downstream Channel List. A value of 0 for Timeout means no time out, i.e. the CM scans all available frequencies once, then moves to the next sub-TLV in the Downstream Channel List.

The CM MUST be capable of processing a Downstream Channel List that contains multiple Default Scanning TLVs.

Type	Length	Value
41.3	2	Timeout

C.1.1.22.4 Examples Illustrating Usage of the Downstream Channel List

Assume that a modem has been provisioned to receive a configuration file with a Downstream Channel List consisting of several single downstream channel (TLV 41.1) entries, a downstream frequency range (TLV 41.2) entry, a default scanning (TLV 41.3) entry and no timeout entries.

When the CM first boots up, it locks onto the first Primary Downstream Channel it can find and goes through initial ranging. After completing the ranging process, the CM downloads the configuration file with the Downstream Channel List. The CM then checks its current Primary Downstream Channel frequency against the frequencies explicitly listed in the single downstream channel (TLV 41.1) entries and the downstream frequency range entry (TLV 41.2) of the Downstream Channel List, ignoring the default scan (TLV 41.3) entry at this point. If the current Primary Downstream Channel is not explicitly in the single downstream channel entries in the list or within the downstream frequency range entry in the list, the CM moves to the first sub-TLV in the TLV 41 list and attempts to lock onto that channel. If the CM is able to lock onto that frequency and the channel is a suitable Primary Downstream Channel, it again tries to range and download a configuration file. Assuming that the CM receives the same configuration file, the CM would then proceed with registration.

If the CM is not able to lock on the first sub-TLV in the Downstream Channel List or the channel is unsuitable for a Primary Downstream Channel, it moves onto the next entry in the list and so on. If the CM reaches the downstream frequency range TLV, it will begin scanning at the downstream frequency range start, updating the frequency by the downstream frequency step size and ending at the downstream frequency range end. If the CM finds a valid Primary Downstream Channel within the downstream frequency range, the CM ranges and downloads a configuration file. Assuming that the configuration file has not changed, the CM continues with registration on that channel.

However, if the CM reaches the default scanning sub-TLV without successfully registering, the CM starts its "default scan" process. If during the course of its default scan, the CM finds a Primary Downstream Channel that it can lock onto, is able to complete ranging and is able to download a configuration file, it will do so. However, at that point, the CM once again checks that the current Primary Downstream Channel is explicitly listed in the Downstream Channel List and acts accordingly.

As a second, less likely example, assume that a CM has been provisioned to receive a configuration file with a Downstream Channel List containing only a default scanning (TLV 41.3) entry. When the CM first boots up, it locks onto the first Primary Downstream Channel it can find and goes through initial ranging. After completing the ranging process, the CM downloads the configuration file with the Downstream Channel List. Since the default scanning is the only parameter in the Downstream Channel List, the current Primary Downstream Channel frequency on which the CM locked is not explicitly included so the CM continues to scan according to its algorithm. The CM will not register on a channel until it receives a configuration file with a downstream frequency explicitly listed in the Downstream Channel List or a configuration file with no Downstream Channel List.

C.1.1.23 Static Multicast MAC Address

The Static Multicast MAC Address TLV configures static multicast MAC addresses for multicast forwarding; the CM behavior based on this TLV is dependant on whether the CM has Multicast DSID Forwarding enabled (as indicated in the modem capabilities encoding of the REG-RSP or REG-RSP-MP). This object may be repeated to configure any number of static multicast MAC addresses. The CM MUST support a minimum of 16 Static Multicast MAC addresses.

If Multicast DSID Forwarding is enabled, the Static Multicast MAC Address TLV informs the CMTS of multicast MAC addresses that need to be labeled with a DSID and communicated to the CM in the REG-RSP or REG-RSP-MP message. The CM MUST NOT forward traffic based on the static multicast MAC address(es) in these encodings when Multicast DSID Forwarding is enabled. When flagged by the Extended CMTS MIC Bitmap, the CM passes this object to the CMTS in REG-REQ or REG-REQ-MP without performing any action. If this TLV is not flagged by the Extended CMTS MIC Bitmap, it will not be forwarded by the CM in the Registration Request and so will have no effect. The CMTS MUST communicate in its REG-RSP or REG-RSP-MP one or more DSIDs for multicast sessions identified by the Static Multicast MAC Address TLV to be forwarded by that CM in this case.

When Multicast DSID Forwarding is disabled, Static Multicast MAC Address TLV configures the CM with a static multicast MAC address that is being provisioned into the CM. The CM MUST forward any multicast frames that match the static multicast MAC address from the cable network to the CMCI subject to the provisions of clause G.4.3 when Multicast DSID Forwarding is Disabled. IGMP has no impact on this forwarding.

When an operator desires to encrypt IP multicast sessions that map to Static Multicast MAC Address TLV the operator must also include Static Multicast Session Encodings in the CM config file. This is because the CMTS controls the encryption based on multicast IP addresses and not based on MAC addresses.

Type	Length	Value
42	6	Static Multicast MAC Address

C.1.1.24 Downstream Unencrypted Traffic (DUT) Filtering Encoding

This parameter enables the CM to perform Downstream Unencrypted Traffic filtering as described in the DOCSIS Layer 2 Virtual Private Network specification [8]. If the CM does not support the DUT Filtering Capability, it MUST ignore the DUT Filtering Encoding TLV.

Type	Length	Value
45	Length/value tuples are specified in [8]	

C.1.1.25 Channel Assignment Configuration Settings

This field is used to convey an assigned Transmit Channel Set and/or Receive Channel Set to be used by a CM via a config file setting which is transmitted to the CMTS in a Registration Request message. It includes two sub-TLVs, one each for transmit and receive channels respectively. There can be multiple instances of each sub-TLV in a single Channel Assignment Configuration Settings encoding, one for each transmit and/or receive channel being assigned to the CM. The list of upstream and/or downstream channels assigned represents the complete list of channels to be assigned to that modem, overriding any other channel assignments that the CMTS might have chosen to make.

If a CMTS receives this field, it MUST either assign only the complete list of assigned transmit and/or receive channels or reject the registration attempt if it is unable to provide all of the assigned channels.

Type	Length	Value
56	N	

C.1.1.25.1 Transmit Channel Assignment Configuration Setting

The US Channel ID to be included in the Transmit Channel Set.

Type	Length	Value
56.1	1	Upstream Channel ID

C.1.1.25.2 Receive Channel Assignment Configuration Setting

The DS Channel Frequency to be included in the Receive Channel Set.

Type	Length	Value
56.2	4	Rx Frequency

C.1.1.26 Upstream Drop Classifier Group ID

The value of this field specifies the list of Upstream Drop Classifier Group IDs [10]. The CMTS uses these Group IDs to instantiate UDCs in the registration response message. The CMTS SHOULD ignore an Upstream Drop Classifier Group ID with a value of zero in the registration request message.

Type	Length	Value
62	n	1 to 255

C.1.1.27 CMTS Static Multicast Session Encoding

The CMTS Static Multicast Session is used by the operator to provide the CMTS with the static ASM or SSM multicast sessions and associated CMIM to which the CM should be configured to forward multicast traffic. To configure static ASM sessions, the CMTS Static Multicast Session Encoding contains the Static Multicast Group Encoding and the Static Multicast CMIM Encoding. To configure static SSM sessions, the CMTS Static Multicast Session Encoding contains the Static Multicast Group Encoding, the Static Multicast Source Encoding and the Static Multicast CMIM Encoding. When flagged by the Extended CMTS MIC Bitmap, the CM passes this object to the CMTS in REG-REQ or REG-REQ-MP without performing any action. If this TLV is not flagged by the Extended CMTS MIC Bitmap, it will not be forwarded by the CM in the Registration Request and so will have no effect.

As described in clause 9.2.4, the CMTS is required to communicate a DSID and associated encodings to the CM in a Registration Response message in response to CMTS Static Multicast Session Encodings present in the Registration Request.

This object may be repeated to configure any number of multicast sessions and associated CMIMs.

Type	Length	Value
64	N	

C.1.1.27.1 Static Multicast Group Encoding

The Static Multicast Group Encoding provides the CMTS with the group address for a multicast session to which the CM will be statically joined. A valid CMTS Static Multicast Session encoding contains exactly one instance of this sub-TLV.

Subtype	Length	Value
64.1	4 (IPv4) or 16 (IPv6)	Multicast group address

C.1.1.27.2 Static Multicast Source Encoding

The Static Multicast Source Encoding provides the CMTS with a source address for a source-specific multicast session to which the CM will be statically joined. A valid CMTS Static Multicast Session encoding may contain multiple instances of this sub-TLV.

Subtype	Length	Value
64.2	4 (IPv4) or 16 (IPv6)	Source IP Address

C.1.1.27.3 Static Multicast CMIM Encoding

The Static Multicast CMIM Encoding provides the CMTS with the CMIM associated with the static multicast session that needs to be communicated to the CM. Each bit of CM interface mask corresponds to a logical or physical interface. (Refer to clause C.1.5.4.4.2.) Multicast CM Interface Mask for details on what interface each bit represents.

A valid CMTS Static Multicast Session encoding contains exactly one instance of this sub-TLV.

Subtype	Length	Value
64.3	N	Static Multicast CMIM

C.1.2 Configuration-File-Specific Settings

The TLVs in the following subclauses are not intended to be forwarded by the CM to the CMTS in the Registration Request message. As such, they are not expected to be included in the E-MIC Bitmap.

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

This is the file name of the software upgrade file for the CM. The file name 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 clause D.1.2. See also clause 12.1.

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 [19] 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.

Enhanced access control for cable modem MIB objects is provided by the SNMPv3 Access View Configuration encoding (see clause C.1.2.14), therefore this object is deprecated. If the configuration file does not contain one or more SNMPv3 Access View Configuration encodings, the CM MAY silently ignore SNMP Write-Access Control encodings. If the configuration file contains one or more SNMPv3 Access View Configuration encodings, the CM MUST silently ignore SNMP Write-Access Control encodings.

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

The value is an SNMP VarBind as defined in [30]. 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:

- The request is treated as fully authorized (it cannot refuse the request for lack of privilege).
- SNMP Write-Control provisions (see clause C.1.2.4) 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 by the CM 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 clause 9.1.2.1). 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 IPv4 TFTP Server

The IPv4 address of the TFTP server on which the software upgrade file for the CM resides. See clauses 12.1 and C.1.2.3.

Type	Length	Value
21	4	TFTP Server's IPv4 Address

C.1.2.8 Software Upgrade IPv6 TFTP Server

The IPv6 address of the TFTP server on which the software upgrade file for the CM resides. See clauses 12.1 and C.1.2.3.

Type	Length	Value
58	16	TFTP Server's IPv6 Address

C.1.2.9 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 the operating mode of the CMs.

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.1.2.9.1 SnmpV3 Kickstart Security Name

Type	Length	Value
34.1	2 to 16	UTF8 Encoded security name

For the ASCII character set, the UTF8 and the ASCII 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.1.2.9.2 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 [47]. 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.1.2.10 Manufacturer Code Verification Certificate

The Manufacturer's Code Verification Certificate (M-CVC) for Secure Software Downloading specified by [15]. The CM config file MUST contain this M-CVC and/or C-CVC defined in clause C.1.2.11 in order to allow the 1.1-compliant CM to download the code file from the TFTP server whether or not the CM is provisioned to run with BPI, BPI+ or none of them. See [15] for details.

Type	Length	Value
32	n	Manufacturer CVC (DER-encoded ASN.1)

If the length of the M-CVC exceeds 254 bytes, the M-CVC is fragmented into two or more successive Type 32 elements. Each fragment, except the last, must be 254 bytes in length. The CM MUST reconstruct 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.1.2.11 Co-signer Code Verification Certificate

The Co-signer's Code Verification Certificate (C-CVC) for Secure Software Downloading specified by [15]. The CM config file MUST contain this C-CVC and/or M-CVC defined in clause C.1.2.10 in order to allow the 1.1-compliant CM to download the code file from TFTP server whether or not the CM is provisioned to run with BPI, BPI+ or none of them. See [15] for details.

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 is fragmented into two or more successive Type 33 elements. Each fragment, except the last, must be 254 bytes in length. The CM MUST reconstruct 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.1.2.12 SNMPv3 Notification Receiver

This TLV specifies a Network Management Station that will receive notifications from the modem when it is in Coexistence mode. Up to 10 of these elements may be included in the configuration file. Please refer to [10] for additional details of its usage.

Type	Length	Value
38	N	Composite

C.1.2.12.1 SNMPv3 Notification Receiver IPv4 Address

This sub-TLV specifies the IPv4 address of the notification receiver.

Type	Length	Value
38.1	4	IPv4 Address

C.1.2.12.2 SNMPv3 Notification Receiver UDP Port Number

This sub-TLV specifies the UDP port number of the notification receiver. If this sub-TLV is not present, the default value of 162 should be used.

Type	Length	Value
38.2	2	UDP port number

C.1.2.12.3 SNMPv3 Notification Receiver Trap Type

Type	Length	Value
38.3	2	Trap type

This sub-TLV specifies the type of trap to send. The trap type may take values:

- 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.

C.1.2.12.4 SNMPv3 Notification Receiver Timeout

This sub-TLV specifies the timeout value to use when sending an Inform message to the notification receiver.

Type	Length	Value
38.4	2	Time in milliseconds

C.1.2.12.5 SNMPv3 Notification Receiver Retries

This sub-TLV specifies the number of times to retry sending an Inform message if an acknowledgement is not received.

Type	Length	Value
38.5	2	Number of retries

C.1.2.12.6 SNMPv3 Notification Receiver Filtering Parameters

This sub-TLV specifies the ASN.1 formatted Object Identifier of the snmpTrapOID value that identifies the notifications to be sent to the notification receiver. SNMP v3 allows the specification of which Trap OIDs are to be sent to a trap receiver. This object specifies the OID of the root of a trap filter sub-tree. All Traps with a Trap OID contained in this trap filter sub-tree MUST be sent by the CM to the trap receiver. This object 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.

Type	Length	Value
38.6	N	Filter OID

C.1.2.12.7 SNMPv3 Notification Receiver Security Name

This sub-TLV specifies the V3 Security Name to use when sending a V3 Notification. This sub-TLV is only used if Trap Type is set to 4 or 5. This name must be a name specified in a config file TLV Type 34 as part of the Diffie-Hellman (DH) Kickstart procedure. The notifications will be sent using the Authentication and Privacy Keys calculated by the modem during the DH Kickstart procedure.

This sub-TLV is not required for Trap Type = 1, 2 or 3 above. If it is not supplied for a Trap type of 4 or 5, then the V3 Notification will be sent in the noAuthNoPriv security level using the security name "@config".

Type	Length	Value
38.7	N	security name

C.1.2.12.8 SNMPv3 Notification Receiver IPv6 Address

This sub-TLV specifies the IPv6 address of the notification receiver.

Type	Length	Value
38.8	16	IPv6 Address

C.1.2.13 SNMPv1v2c Coexistence Configuration

This object specifies the SNMPv1v2c Coexistence Access Control configuration of the CM. This object does not preclude using TLV-11 to configure directly SNMPv3 tables. The CM MUST support a minimum of 10 SNMPv1v2c Coexistence TLVs. This TLV creates entries in SNMPv3 tables as specified in [10].

If the configuration file contains SNMPv1v2 Coexistence Configuration encodings, the CM MUST reject the configuration file if the SNMPv1v2c Community Name and SNMPv1v2c Transport Address Access sub-TLVs are not present. The CM MUST support multiple instances of sub-TLV 53.2 SNMPv1v2c Transport Address Access. The CM MUST reject a config file if a TLV includes repeated sub-TLVs other than sub-TLV 53.2. The CM MUST reject the config file if a CM created entry in a SNMP table is rejected for syntax conflicts or reaches the limit in the number of entries the CM support for that table or the mapped SNMPv3 entry already exist. The CM MUST reject the config file if the TLV has an invalid length or if any of the sub-TLVs have an invalid length or value.

Type	Length	Value
53	N	Composite

NOTE: The number of entries a CM can support in SNMPv3 tables is independent of the number of TLVs the CM needs to support to be processed as SNMP tables entries.

C.1.2.13.1 SNMPv1v2c Community Name

This sub-TLV specifies the Community Name (community string) used in SNMP requests to the CM.

Type	Length	Value
53.1	1..32	Text

C.1.2.13.2 SNMPv1v2c Transport Address Access

This sub-TLV specifies the Transport Address and Transport Address Mask pair used by the CM to grant access to the SNMP entity querying the CM. The CM MUST reject a config file if a sub-TLV Transport Address Access has more than one sub-TLV 53.2.1 or 53.2.2.

Type	Length	Value
53.2	n	Variable

C.1.2.13.2.1 SNMPv1v2c Transport Address

This sub-TLV specifies the Transport Address to use in conjunction with the Transport Address Mask used by the CM to grant access to the SNMP entity querying the CM.

Type	Length	Value
53.2.1	6 or 18	Transport Address

NOTE: Length is 6 bytes for IPv4 and 18 bytes for IPv6. Two additional bytes are added to the IP address length for the port number. Refer to the clause "SNMPv1v2c Coexistence Configuration config file TLV" in [10] for details.

C.1.2.13.2.2 SNMPv1v2c Transport Address Mask

This sub-TLV specifies the Transport Address Mask to use in conjunction with the Transport Address used by the CM to grant access to the SNMP entity querying the CM. This sub-TLV is optional.

Type	Length	Value
53.2.2	6 or 18	Transport Address Mask

NOTE: Length is 6 bytes for IPv4 and 18 bytes for IPv6. Two additional bytes are added to the IP address length for the port number. Refer to the clause "SNMPv1v2c Coexistence Configuration config file TLV" in [10] for details.

C.1.2.13.3 SNMPv1v2c Access View Type

This sub-TLV specifies the type of access to grant to the community name of this TLV. Sub-TLV 53.3 is optional. If sub-TLV 53.3 is not present in TLV 53, the default value of the access type to grant to the community name specified in sub-TLV 53.1 is Read-only.

Type	Length	Value
53.3	1	1 Read-only 2 Read-write

C.1.2.13.4 SNMPv1v2c Access View Name

This sub-TLV specifies the name of the view that provides the access indicated in sub-TLV SNMPv1v2c Access View Type.

Type	Length	Value
53.4	1..32	String

C.1.2.14 SNMPv3 Access View Configuration

This object specifies the SNMPv3 Simplified Access View configuration of the CM. This object does not preclude using TLV-11 to configure directly SNMPv3 tables. This TLV creates entries in SNMPv3 tables as specified in [10].

The CM MUST reject the config file if the SNMPv3 Access View Configuration encoding is present but the SNMPv3 Access View Name sub-TLV is not present. The CM MUST support multiple TLVs with the same SNMPv3 Access View Name TLV. The CM MUST reject the config file if more than one sub-TLV is included in a TLV. The CM MUST reject the config file if a CM created entry in a SNMP table is rejected for Syntax conflicts or reaches the limit in the number of entries the CM support for that table or the mapped SNMPv3 entry already exist. The CM MUST reject the config file if the TLV has an invalid length or if any of the sub-TLVs have an invalid length or value.

Type	Length	Value
54	N	Composite

NOTE: The number of entries a CM can support in SNMPv3 tables is independent of the number of TLVs the CM needs to support to be processed as SNMP tables entries.

C.1.2.14.1 SNMPv3 Access View Name

This sub-TLV specifies the administrative name of the View defined by this TLV.

Type	Length	Value
54.1	1..32	Text

C.1.2.14.2 SNMPv3 Access View Subtree

This sub-TLV specifies an ASN.1 formatted object Identifier that represents the filter sub-tree included in the Access View TLV. The CM MUST accept only encoded values that start with the ASN.1 Universal type 6 (Object Identifier) byte, then the ASN.1 length field, then the ASN.1 encoded object identifier components. For example the sub-tree 1.3.6 is encoded as 0x06 0x02 0x2B 0x06. If this sub-TLV is not included in the TLV the CM MUST use as default the OID sub-tree 1.3.6.

Type	Length	Value
54.2	N	OID

The CM MUST assign default OID value 1.3.6 to SNMPv3 Access View Subtree if TLV 54 is present but sub-TLV 54.2 is not included.

C.1.2.14.3 SNMPv3 Access View Mask

This sub-TLV specifies the bit mask to apply to the Access View Subtree of the Access View TLV.

Type	Length	Value
54.3	0..16	Bits

The CM MUST assign a zero-length string to SNMPv3 Access View Mask TLV 54.3 if TLV 54 is present but this sub-TLV is not included.

C.1.2.14.4 SNMPv3 Access View Type

This sub-TLV specifies the inclusion or exclusion of the sub-tree indicated by SNMPv3 Access View Subtree sub-TLV 54.2 in the SNMPv3 Access View Configuration TLV 54. The value 1 indicates the sub-tree of SNMPv3 Access View SubTree is included in the Access View. The value 2 indicates the sub-tree of SNMPv3 Access View Sub Tree is excluded from the Access View.

Type	Length	Value
54.4	1	1 included 2 excluded

The CM MUST assign the value included to SNMPv3 Access View Type sub-TLV 54.4 if TLV 54 is present but this sub-TLV is not included.

C.1.2.15 SNMP CPE Access Control

If the value of this field is a 1, the CM MUST allow SNMP access from any CPE attached to it. If the value of this field is a 0, the CM MUST NOT allow SNMP Access from any CPE attached to it.

Type	Length	Value
55	1	0 Disable 1 Enable

The CM MUST disable SNMP access from CPEs connected to the cable modem unless this TLV is present in the config file with value equal to 1.

C.1.2.16 Management Event Control Encoding

This TLV specifies the mechanism to individually enable DOCSIS events. The CM MUST support one or more instances of TLV 66 in the config file. The CM MUST ignore TLV 66 instances containing the same EventID value in the config file.

Type	Length	Value
66	4	32-bit Event ID or 0 See [10]

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.

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: 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.

The CM MUST include all these capabilities in both the Registration Request and option 125 of the DHCP request unless the description of the capability explicitly prohibits this (such as for capabilities that are not subject to negotiation). The CMTS MUST include Modem Capabilities in the Registration Response as indicated in clause 6.4.8.

C.1.3.1.1 Concatenation Support

If the value field is a 1 the CM requests pre-3.0 DOCSIS concatenation support from the CMTS.

Type	Length	Value
5.1	1	1 or 0

If the value field in REG-RSP or REG-RSP-MP is 0, the CM MUST disable concatenation.

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: DOCSIS v2.0
		3: DOCSIS v3.0
		4-255: 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 clause G.2). This capability is provided by the CM for the benefit of the CMTS, the operation of the CM is not affected by the value returned by the CMTS.

C.1.3.1.3 Fragmentation Support

If the value field is a 1 the CM requests pre-3.0 DOCSIS 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 not in the Registration Request. If a CMTS does receive this capability within 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 indicates the BPI support of the CM.

Type	Length	Value
5.6	1	0: BPI Support
		1: BPI Plus Support
		2 to 255: Reserved

If the value field in REG-RSP or REG-RSP-MP is 0, the CM MUST NOT operate in BPI Plus mode.

C.1.3.1.7 Downstream SAID Support

This field shows the number of Downstream SAIDs that the CM can support.

Type	Length	Value
5.7	1	Number of Downstream SAIDs that the CM can support

If the number of Downstream SAIDs is 0, the Modem can support only one Downstream SAID.

C.1.3.1.8 Upstream Service Flow Support

This field shows the number of Upstream Service Flows that the CM supports which can be used for any Service Flow Scheduling Type.

Type	Length	Value
5.8	1	Number of Upstream Service Flows of all types the CM can support

If the number of Upstream Service Flows is 0, the CM can support only one Upstream Service Flow.

NOTE: In pre-3.0 DOCSIS specifications, this capability was referred to as "Upstream SID Support". Since the number of Upstream SIDs is equivalent to the number of Upstream Service Flows in pre-3.0 DOCSIS, the revisions to this capability are fully backward compatible.

C.1.3.1.9 Optional Filtering Support

This field shows the optional filtering support in the CM. Bits are set to 1 to indicate that support for optional filtering.

Type	Length	Value
5.9	1	Packet Filtering Support Bitmap
		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 not in the Registration Request. If a CMTS does receive this capability within a Registration Request, it MUST return the capability with the same value in the Registration Response.

C.1.3.1.10 Transmit Pre-Equalizer Taps per Modulation Interval

This field shows the maximal number of pre-equalizer taps per modulation interval T supported by the CM. The CM MUST include this capability in the Registration Request with the value 1.

NOTE: All CMs support, at a minimum, T-spaced equalizer coefficients. Support of 2 or 4 taps per modulation interval was optional for DOCSIS 1.0 and 1.1 CMs, while DOCSIS 2.0 and 3.0 CMs are required to only support 1 tap per modulation interval. If this tuple is missing, it is implied that the CM only supports T 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. The CM MUST include this capability in the Registration Request with value 24.

NOTE: All CMs support an equalizer length of at least 8 symbols. Support of up to 64 T-spaced, T/2-spaced or T/4-spaced taps was optional for DOCSIS 1.0 and 1.1 CMs, while DOCSIS 2.0 and 3.0 CMs are required to support exactly 24 taps. 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

This field indicates the DCC support of the CM.

Type	Length	Value
5.12	1	0 = DCC is not supported
		1 = DCC is supported

C.1.3.1.13 IP Filters Support

This field shows the number of IP filters that are supported by the CM.

Type	Length	Value
5.13	2	64 to 65535

NOTE: This CM capability is not subject to negotiation with the CMTS.

The CM MUST include this capability in the DHCP request, but not in the Registration Request. If a CMTS does receive this capability within a Registration Request, it MUST return the capability with the same value in the Registration Response.

C.1.3.1.14 LLC Filters Support

This field shows the number of LLC filters that are supported by the CM.

Type	Length	Value
5.14	2	10-65535

NOTE: This CM capability is not subject to negotiation with the CMTS.

The CM MUST include this capability in the DHCP request, but not in the Registration Request. If a CMTS does receive this capability within a Registration Request, it MUST return the capability with the same value in the Registration Response.

C.1.3.1.15 Expanded Unicast SID Space

This field indicates if the CM can support the expanded unicast SID space.

Type	Length	Value
5.15	1	0 = Expanded Unicast SID space is not supported
		1 = Expanded Unicast SID space is supported

C.1.3.1.16 Ranging Hold-Off Support

The CM indicates support for the Ranging Hold-Off Feature by reporting its Ranging Class ID in the value field. The low order 16 bits of the Ranging Class ID are comprised of a static bit map which indicates the device type. The CM sets the bits of the devices to 1 in the bit map. Only a stand-alone CM will set Bit#0. For example, a standalone CM would report a value of 1; a CM with a CableHome PS or an eRouter would report a value of 2; a CM with a PacketCable MTA and an ePS would report a value of 6; an eSTB would report a value of 8 although it contained an eCM. Bits 16 thru 31 are derived from the Configuration File as described in clause C.1.1.18.1.4. The Ranging Class ID is not negotiable. The CM MUST ignore the value field in the REG-RSP or REG-RSP-MP.

Type	Length	Value
5.16	4	Ranging Class ID (bitmap)
		Bit #0: CM
		Bit #1: ePS or eRouter
		Bit #2: eMTA or EDVA
		Bit #3: DSG/eSTB
		Bits 4 through 15: Reserved
		Bits 16 through 31: CM Ranging Class ID Extension

C.1.3.1.17 L2VPN Capability

This capability indicates whether the CM is compliant with the DOCSIS Layer 2 Virtual Private Network feature as defined in [8]. The CM MAY support the DOCSIS Layer 2 Virtual Private Network feature as defined in [8].

Type	Length	Value
5.17	Length/value tuples are specified in [8]	

C.1.3.1.18 L2VPN eSAFE Host Capability

This capability encoding informs the CMTS of the type and MAC address of an eSAFE host embedded with a CM that supports the L2VPN feature. A CM MUST NOT include L2VPN eSAFE Host Capability TLV in the Registration Request or DHCP Option 60 if it does not indicate support for [8] via the L2VPN Capability encoding or if it is not embedded with any eSAFE host.

Type	Length	Value
5.18	Length/value tuples are specified in [8]	

C.1.3.1.19 Downstream Unencrypted Traffic (DUT) Filtering

This capability indicates whether the CM supports the DUT Filtering feature as defined in the DOCSIS Layer 2 Virtual Private Network specification [8]. The CM MAY support DUT Filtering. A CM MUST NOT include the Downstream Unencrypted Traffic (DUT) Filtering TLV in the Registration Request or DHCP Option 60 if it does not indicate support for [8] via the L2VPN Capability encoding.

Type	Length	Value
5.19	Length/value tuples are specified in [8]	

C.1.3.1.20 Upstream Frequency Range Support

This field shows the upstream frequency range(s) supported by the CM. This setting is independent of the upstream frequency range that is configured in the MDD.

Type	Length	Value
5.20	1	0: Standard Upstream Frequency Range (see [12])
		1: Standard Upstream Frequency Range and Extended Upstream Frequency Range (see [12])
		2 to 255: Reserved

NOTE: If this CM capability setting is not included, the CM is capable only of the Standard Upstream Frequency Range.

C.1.3.1.21 Upstream Symbol Rate Support

This field indicates whether the CM is able to support various upstream symbol rates. CMs are required to support the 5 120 ksps, 2 560 ksps and 1 280 ksps rates [12].

Bit #0 is the LSB of the Value field. Bits are set to 1 to indicate support of the particular symbol rate.

Type	Length	Value
5.21	1	Bit #0 = 160 ksps symbol rate supported
		Bit #1 = 320 ksps symbol rate supported
		Bit #2 = 640 ksps symbol rate supported
		Bit #3 = 1 280 ksps symbol rate supported
		Bit #4 = 2 560 ksps symbol rate supported
		Bit #5 = 5 120 ksps symbol rate supported
		All other bits are reserved

If this encoding is not included, it is assumed that the CM supports 5 120 kps, 2 560 kps, 1 280 kps, 640 kps, 320 kps and 160 kps symbol rates.

C.1.3.1.22 Selectable Active Code Mode 2 Support

This field indicates whether the CM supports Selectable Active Code (SAC) Mode 2.

Type	Length	Value
5.22	1	0: SAC Mode 2 is not supported
		1: SAC Mode 2 is supported
		2 to 255: Reserved

NOTE: If this CM capability setting is not included, the CM is assumed to be not capable of supporting SAC Mode 2.

C.1.3.1.23 Code Hopping Mode 2 Support

This field indicates whether the CM supports Code Hopping Mode 2.

Type	Length	Value
5.23	1	0: Code Hopping Mode 2 is not supported
		1: Code Hopping Mode 2 is supported
		2 to 255: Reserved

NOTE: If this CM capability setting is not included, the CM is assumed to be not capable of supporting Code Hopping Mode 2.

C.1.3.1.24 Multiple Transmit Channel Support

This field shows the number of upstream transmitters that the CM can support. A non-zero value indicates that this CM supports the Multiple Transmit Channel Mode of operation, which includes:

- Continuous Concatenation and Fragmentation (CCF), including the use of Segment headers.
- Queue-depth based bandwidth requests.
- Multiple requests outstanding.
- SID Clusters.
- T4 Timeout Multiplier.
- Use of assigned burst profile corresponding to IUC in the grant.
- One-fill of FEC code words, as opposed to zero-fill.
- Other new request and transmission rules.

Type	Length	Value
5.24	1	Number of upstream transmitters that the CM can support

Since, for 3.0 operation, only the three highest symbol rates are operative, this number is equivalent to the number of 1,28 Msps transmitters that the CM can support. If this CM capability setting is not included or the number of upstream transmitters is 0, the CM does not support Multiple Transmit Channel Mode. The CM MUST indicate support for 4 or more upstream transmitters.

If the CMTS returns a value of 0 in the REG-RSP or REG-RSP-MP, the CM MUST disable Multiple Transmit Channel Mode as described above.

C.1.3.1.25 5,12 Msps Upstream Transmit Channel Support

This field shows the maximum number of upstream channels at a symbol rate of 5,12 Msps that the CM can support.

Type	Length	Value
5.25	1	Number of upstream channels at 5,12 Msps that the CM can support

If this CM capability setting is not included or the number of upstream channels is 0, the CM can support only one upstream channel at 5,12 Msps. A CM that can support N channels at symbol rate 5,12 Msps can support N channels at equal or lower symbol rates.

C.1.3.1.26 2,56 Msps Upstream Transmit Channel Support

This field shows the maximum number of upstream channels at symbol rate 2,56 Msps that the CM can support.

Type	Length	Value
5.26	1	Number of upstream channels at 2,56 Msps that the CM can support

If this CM capability setting is not included or the number of upstream channels is 0, the CM can support only one upstream channel at 2,56 Msps. A CM that can support N channels at symbol rate 2,56 Msps can support N channels at equal or lower symbol rates.

C.1.3.1.27 Total SID Cluster Support

This field shows the total number of SID Clusters that the CM can support.

Type	Length	Value
5.27	1	Total number of SID Clusters supported

The CM MUST support a total number of SID Clusters at least two times the number of Upstream Service Flows supported as reported in clause C.1.3.1.8 plus one SID Cluster for the number of UGS or UGS-AD only Service Flows as reported in clause C.1.3.1.36.

C.1.3.1.28 SID Clusters per Service Flow Support

This field shows the maximum number of SID Clusters that can be assigned to a service flow for this CM.

Type	Length	Value
5.28	1	2 to 8 Maximum number of SID Clusters per Service Flow

C.1.3.1.29 Multiple Receive Channel Support

This value is used by the CM to indicate that it can receive more than one downstream channel simultaneously. This encoding gives the maximum number of separately identified Receive Channels (RCs) that the CM can support. A non-zero value indicates that this CM:

- supports DSID encodings in the REG-RSP-MP and DBC-REQ;
- supports transmitting its RCPs in the REG-REQ-MP;
- supports receiving an RCC encoding in the REG-RSP-MP and DBC-REQ.

Type	Length	Value
5.29	1	Maximum number N of physical downstream Receive Channels identified on the CM. RCs are identified within the CM with an RCID from 1 to N

If the CMTS returns a value of 0 in the REG-RSP or REG-RSP-MP, the CM MUST disable Multiple Receive Channel Support as described above. If the CM omits this encoding or the CMTS returns a value of 0 for this encoding, then the CMTS MUST NOT return a non-zero value for the Multicast DSID Forwarding Capability encoding (clause C.1.3.1.33) and the Multiple Transmit Channel Support encoding (clause C.1.3.1.24) in its REG-RSP or REG-RSP-MP message to the CM.

C.1.3.1.30 Total Downstream Service ID (DSID) Support

The value of this field indicates the maximum total number of Downstream Service IDs (DSIDs) that the CM can recognize for filtering purposes.

Type	Length	Value
5.30	1	32 to 255

C.1.3.1.31 Resequencing Downstream Service ID (DSID) Support

The value of this field indicates the number of resequencing DSIDs (resequencing contexts) that the CM can support simultaneously. This number must be no higher than the maximum number of DSIDs supported (see clause C.1.3.1.30).

Type	Length	Value
5.31	1	16 to 255

C.1.3.1.32 Multicast Downstream Service ID (DSID) Support

The value of this field indicates the number of multicast Downstream Service IDs (DSIDs) used by the CMTS to label multicast streams that the CM can support simultaneously. The value of this field also indicates the number of multicast DSIDs on which the CM supports multicast PHS Rules. This number MUST be no higher than the Total DSID Support (see clause C.1.3.1.30).

Type	Length	Value
5.32	1	16 to 255

C.1.3.1.33 Multicast DSID Forwarding

The value is used by the CM to indicate its level of support for multicast DSID forwarding that is introduced in DOCSIS 3.0. The CM reports one of three levels of support for Multicast DSID Forwarding.

- **No support for multicast DSID forwarding (0):** A CM reports this value if it cannot forward multicast traffic based on the DSID. A CM that reports this value cannot perform DSID-indexed PHS.
- **GMAC Explicit Multicast DSID Forwarding (1):** A CM reports this value if it is capable of forwarding multicast traffic labeled with a known DSID but is requesting an explicit list of destination GMAC addresses. A CM that reports this value is capable of performing DSID-indexed PHS.
- **GMAC Promiscuous Multicast DSID Forwarding (2):** A CM reports this value if it is capable of forwarding multicast traffic based only on the DSID. A CM that reports this value is capable of performing DSID-indexed PHS.

Since a CM that reports support for either type of multicast DSID forwarding, GMAC explicit or GMAC promiscuous, forwards all downstream multicast traffic based on the DSID, a CM is considered to be capable of Multicast DSID forwarding if it reports a value of 1 or 2.

The CM MUST indicate support for GMAC Promiscuous Multicast DSID Forwarding.

A CMTS that returns a non-zero value of the Multicast DSID Forwarding Support capability encoding to a CM in a REG-RSP or REG-RSP-MP is said to "enable" Multicast DSID Forwarding at the CM.

If a CM reports that it is capable of Multicast DSID Forwarding with the value of 1 or 2, the CMTS MAY return a value of 0 for the encoding in its REG-RSP or REG-RSP-MP in order to "disable" Multicast DSID Forwarding for a CM. If the CMTS returns a value of 0 in the REG-RSP or REG-RSP-MP, the CM MUST disable its Multicast DSID Forwarding.

The CMTS MUST NOT return a value of 1 for the Multicast DSID Forwarding Capability encoding in its REG-RSP or REG-RSP-MP message to the CM unless the CM advertised a capability of 1. If the CM advertises a capability of 1, the CMTS has the option of returning a value of 2 (see clause G.4.2.2).

Type	Length	Value
5.33	1	0 = No support for multicast DSID forwarding 1 = Support for GMAC explicit multicast DSID forwarding 2 = Support for GMAC promiscuous multicast DSID forwarding 3 to 255 = Reserved

C.1.3.1.34 Frame Control Type Forwarding Capability

This value is used by the CM to indicate support for forwarding traffic with the Isolation PDU MAC Header (the FC_Type field with a value of 10).

Type	Length	Value
5.34	1	0 = Isolation Packet PDU MAC Header (FC_Type of 10) is not forwarded 1 = Isolation Packet PDU MAC Header (FC_Type of 10) is forwarded 2 to 255 = Reserved

The CM MUST indicate support for forwarding traffic with the FC_Type field set to a value of 10.

C.1.3.1.35 DPV Capability

This value is used by the CM to indicate support for the DOCSIS Path Verify Feature.

Type	Length	Value
5.35	1	Bit 0: U1 supported as a Start Reference Point for DPV per Path. Bit 1: U1 supported as a Start Reference Point for DPV per Packet. Bits 2 to 7 are reserved

C.1.3.1.36 Unsolicited Grant Service/Upstream Service Flow Support

This field shows the number of additional Service Flows that the CM supports which can be used only for Unsolicited Grant Service. This includes UGS and UGS/AD scheduling services.

Type	Length	Value
5.36	1	Number of additional service flows that the CM can support which can be used only for Unsolicited Grant Service Flows

C.1.3.1.37 MAP and UCD Receipt Support

This field indicates whether or not the CM can support the receipt of MAPs and UCDs on any downstream channel or if it can only receive MAPs and UCDs on the Primary Downstream Channel.

Type	Length	Value
5.37	1	0 = CM cannot support the receipt of MAPs and UCDs on downstreams other than the Primary Downstream Channel
		1 = CM can support the receipt of MAPs and UCDs on downstreams other than the Primary Downstream Channel

The CM MUST support a capability of 1 (CM can support the receipt of MAPs and UCDs on any downstream channel). If the CMTS sets this capability to 0 in the REG-RSP or REG-RSP-MP, the CM MUST look for MAPs and UCDs only on the Primary Downstream Channel.

If the CMTS receives a REG-REQ or REG-REQ-MP message with the MAP and UCD Receipt Support modem capability of 0, then it MUST provide MAPs and UCDs for that CM on its Primary Downstream Channel.

C.1.3.1.38 Upstream Drop Classifier Support

This field shows the number of Upstream Drop Classifiers that are supported by the CM.

Type	Length	Value
5.38	2	64 to 65535

The CM MUST indicate support for at least 64 Upstream Drop Classifiers.

NOTE: The number of Upstream Drop Classifiers supported by the CM is not subject to negotiation with the CMTS. The CM MUST include this capability in both the DHCP request and the Registration Request. The CMTS disables Upstream Drop Classification by returning a value of zero for the Upstream Drop Classifier Support in the registration response. The CM handling of the CMTS response is described in clause 7.5.1.2.2.

C.1.3.1.39 IPv6 Support

This value is used by the CM to indicate support for IPv6 provisioning and management.

Type	Length	Value
5.39	1	0 = IPv6 is not supported 1 = IPv6 is supported 2 to 255 = Reserved

The CM MUST indicate support for IPv6.

C.1.3.1.40 Extended Upstream Transmit Power Capability

The Extended Upstream Transmit Power Capability is used to communicate the CM's support for increasing P_{\max} values as described in [12] for upstream channels when MTC Mode is enabled. If the default value for P_{\max} for an individual upstream channel is lower than the capability encoding, the CM and CMTS adjust the value to be equal to the capability encoding. If the default value for P_{\max} is the same as or higher than the capability encoding the CM and CMTS retain the default value. Note that this capability only affects the value of P_{\max} ; the CMTS controls the CM's transmit power via the Dynamic Range Window (see clause 8.3.3 and [12] for more details).

The CM MUST report the Extended Upstream Transmit Power Capability in units of one-quarter dB. A CM capability value of zero indicates that the CM does not support an extension to its Upstream Transmit Power. If the CMTS returns a non-zero value that is different from the value the CM sent in the REG-REQ-MP, this indicates that the CM and CMTS are not synchronized and the CM MUST re-initialize the MAC with Initialization Reason "REG_RSP_NOT_OK" (7). If the CMTS returns a zero value or if the TLV is absent, the CM does not extend its upstream transmit power as defined in [12].

The CMTS MUST either confirm the CM's capability by responding with the same value communicated by the CM or disable the Extended Upstream Transmit Power capability by responding with a value of zero. By default, the CMTS MUST confirm the value, unless a mechanism is provided to administratively configure this setting on and off.

Type	Length	Value
5.40	1	0, 205 to 244 (units of one-quarter dB)

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. The Vendor ID is not used as a stand-alone configuration file element. The Vendor ID **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 IPv4 Address' (see clause C.1.1.9).

Type	Length	Value
12	4	IPv4 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

The Class ID is the class-of-service class from the request which is not available.

The Type is the specific class-of-service object within the class which caused the request to be rejected.

The Confirmation Code is defined in clause C.4.

C.1.3.5 Vendor-Specific Capabilities

Vendor-specific data about the CM, that is to be included in the REG-REQ or REG-REQ-MP, but which is not part of the configuration file, if present, **MUST** be encoded in the vendor specific capabilities (VSC) (code 44) using the Vendor ID field (refer to clause C.1.3.2) to specify which TLV tuples apply to which vendors products. The Vendor ID **MUST** be the first TLV embedded inside VSC. 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. There **MUST NOT** be more than one Vendor ID TLV inside a single VSC.

Type	Length	Value
44	n	Per vendor definition

EXAMPLE:

- Configuration with vendor A specific fields and vendor B specific fields:
 - VSC (44) + n (number of bytes inside this VSC).
 - 8 (Vendor ID Type) + 3 (length field) + Vendor ID of Vendor.

- Vendor Specific Type #1 + length of the field + Value #1.
- Vendor Specific Type #2 + length of the field + Value #2

C.1.3.6 CM Initialization Reason

For debugging and system maintenance it is useful to know what caused a CM to initialize. When a CM performs a MAC initialization it has to retain the Initialization Reason. After initialization the CM will attempt to come online. When it sends a REG-REQ or REG-REQ-MP it reports the Initialization Reason in the REG-REQ or REG-REQ-MP using the "CM Initialization Reason" TLV. The CM MUST include this TLV in the REG-REQ or REG-REQ-MP.

Type	Length	Value
57	1	Initialization Code

Table C-3 outlines the initialization reasons and the associated Initialization Codes.

Table C-3: Initialization Reasons and Codes

Initialization Reason	Initialization Code
POWER-ON	1
T17_LOST-SYNC	2
ALL_US_FAILED	3
BAD_DHCP_ACK	4
LINK_LOCAL_ADDRESS_IN_USE	5
T6_EXPIRED	6
REG_RSP_NOT_OK	7
BAD_RCC_TCC	8
FAILED_PRIM_DS	9
TCS_FAILED_ON_ALL_US	10
MTCM_CHANGE	15
T4_EXPIRED	16
NO_PRIM_SF_USCHAN	17
CM_CTRL_INIT	18
DYNAMIC-RANGE-WINDOW-VIOLATION	19

C.1.4 Dynamic-Message-Specific Encodings

These encodings are not found in the configuration file, nor in the Registration Request/Response signalling. They are only found in DSA-REQ, DSA-RSP, DSA-ACK, DSC-REQ, DSC-RSP, DSC-ACK, DSD-REQ, DBC-REQ, DBC-RSP and DBC-ACK messages (clause 6.4.12 through clause 6.4.18 and clause 6.4.29 through clause 6.4.31).

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. For Dynamic Messages, other than the DBC-REQ message, the message digest MUST be performed over all of the Dynamic Message parameters starting immediately after the MAC Management Message Header and up to, but not including the HMAC Digest setting, in the order in which they appear within the packet. For the DBC-REQ Message, the message digest MUST be performed over all the TLV Encoded Parameters (i.e. not including fixed fields such as the Number of Fragments and Fragment Sequence Number) up to, but not including, the HMAC-Digest setting, in the order in which they appear within the re-assembled 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 [15].

This parameter contains a keyed hash used for message authentication. The HMAC algorithm is defined in [36]. 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 [63].

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 the present document, but include [24].

The Authorization Block MAY be present in CM-initiated DSA-REQ and DSC-REQ messages and CMTS-initiated DSA-RSP and DSC-RSP messages. This parameter MUST NOT be present in CMTS-initiated DSA-REQ and DSC-REQ messages, 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

This value shows the key sequence number of the [15] 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.1.5 Registration, Dynamic Service and Dynamic Bonding Settings

These encodings report the physical capabilities and configuration of downstream receive channels and upstream transmit channels on CMs capable of multiple channel operation.

C.1.5.1 Transmit Channel Configuration (TCC)

This field defines operations to be performed on an upstream channel in the Transmit Channel Set. It can be used in the Registration and DBC MAC Management Messages. If the CMTS confirms a Multiple Transmit Channel Support TLV with a value greater than zero, the CMTS is required to include the TCC TLV in the REG-RSP-MP. If the CMTS enables Multiple Receive Channel mode and sets the Multiple Transmit Channel Support TLV to zero, either by confirming a CM capability of zero or by disabling Multiple Transmit Channel Support for a modem which indicated support via a non-zero value, the CMTS is permitted to include the TCC TLV in the REG-RSP-MP (clause 10.2.6.2). If the CMTS includes the TCC TLV in the REG-RSP-MP, then it uses DBC messaging (as opposed to DCC or UCC messaging) to change the CM's upstream channel(s). If the CMTS does not include the TCC TLV in the REG-RSP-MP, then it does not use DBC messaging to change the CM's upstream channel(s); instead, it uses DCC or UCC messaging for this purpose.

The value field of this TLV contains a series of sub-types.

Type	Length	Value
46	N	

The CMTS MAY include this TLV multiple times within a single message. If the length of the Transmit Channel Configuration (TCC) exceeds 254 bytes, the TCC MUST be fragmented into two or more successive Type 46 elements. Each subsequent TCC fragment MUST begin with a sub-TLV which always contains a complete sub-TLV value unless specified otherwise in the description of the sub-TLV, in which case it could contain a sub-set of the octets of that sub-TLV (e.g. see clause C.1.5.1.5). In other words, a sub-TLV instance value cannot span Type 46 TLV fragments without the Type-Length encoding corresponding to that sub-TLV. If it fragments the TCC Encoding, the CMTS MUST ensure that the fragments arrive in order at the CM, as the CM is not required to resequence out-of-order TCC Encoding fragments.

C.1.5.1.1 Transmit Channel Configuration (TCC) Reference

The CMTS MUST assign a unique Transmit Channel Configuration (TCC) Reference per TCC (Type 46 TLV). The CMTS MUST encode this TLV as the first TLV in any complete Type 46 encoding. In a fragmented TCC encoding, the CMTS MUST encode the TCC Reference as the first TLV in the first fragment. In a fragmented TCC encoding, the CMTS MAY also encode the TCC Reference as the first TLV in subsequent fragments. If it encodes the TCC Reference as the first TLV in subsequent fragments of a TCC encoding, the CMTS MUST use the TCC Reference value encoded in the first fragment.

When it receives a fragmented TCC encoding, the CM MUST NOT consider the TCC encoding invalid if the TCC Reference is the first TLV in only the first fragment or if the TCC Reference is the first TLV in each of the fragments.

Type	Length	Value
46.1	1	0-255: TCC Reference ID

C.1.5.1.2 Upstream Channel Action

The value of this field is used by the CMTS to inform the CM of the action to be performed. These actions include adding the upstream channel to the Transmit Channel Set, changing the ranging SID associated with an upstream channel in the Transmit Channel Set, deleting the upstream channel from the Transmit Channel Set or replacing the upstream channel within the Transmit Channel Set with a new channel.

A value of "Change" (2) is used to change the ranging SID associated with the upstream channel in the Transmit Channel Set or to change the value of the Dynamic Range Window.

A value of "Re-range" (5) is used to re-range all upstream channels that are included in both the old and the new TCC (any channels not being added, deleted or replaced) according to the initialization technique provided (refer to clause C.1.5.1.7). The CM does not re-range upstream channels which are being added, deleted or replaced. This action is required when the primary downstream channel is being changed or affected by implicit or explicit changes in the Receive Module.

A value of "No Action" (0) is provided to allow the TCC to be included in a message even when the specified Upstream Channel ID is already in use by the CM. This action indicates that no changes are required for the CM to continue using the upstream channel. The CMTS MUST NOT include a TCC Encoding with an Upstream Channel Action of "No Action" in the DBC-REQ message if the DBC-REQ message includes a TCC Encoding with an Upstream Channel Action of "Re-Range".

This TLV MUST be included exactly once in the TCC.

Type	Length	Value
46.2	1	0 = No Action 1 = Add 2 = Change 3 = Delete 4 = Replace 5 = Re-range 6 to 255: Reserved

C.1.5.1.3 Upstream Channel ID

This TLV **MUST** be included exactly once in each TCC. It is the ID of the Upstream Channel being operated on. When the action is Replace (4), this ID is the channel being replaced.

When the action is Re-range (5), the value of the upstream channel **MUST** be 0.

When the Dynamic Range Window TLV is included in the TCC, the value of the upstream channel ID **MUST** be 0.

Type	Length	Value
46.3	1	0 = All upstream channels (used with an upstream channel action of Re-Range or inclusion of Dynamic Range Window TLV in the TCC) 1 to 255 = Upstream Channel ID

C.1.5.1.4 New Upstream Channel ID

When the Upstream Channel Action is Replace (4), this TLV **MUST** be included exactly once in the TCC. It **MUST NOT** be present for any other Upstream Channel Action values. This TLV contains the Upstream Channel ID of the new channel which is replacing an existing channel.

Type	Length	Value
46.4	1	1 to 255: Upstream Channel ID

C.1.5.1.5 UCD

The CMTS includes this TLV when the Upstream Channel Action is either Add or Replace so that the CM will not have to wait for a UCD message for the new upstream channel. Including the UCD in the TCC encoding allows the CM to validate the Dynamic Range Window for the commanded TCS prior to making any changes.

The CMTS **MUST** include the UCD encoding within a TCC when the Upstream Channel Action is Add or Replace. The CMTS **MUST NOT** include the UCD encoding within a TCC when the Upstream Channel Action is No action, Change, Delete or Re-range. The CM **MUST** observe the UCD encoding.

Type	Length	Value
46.5	N	

This TLV includes all parameters for the UCD message as described in clause 6.4.3, except for the MAC Management Header. The CMTS **MUST** ensure that the change count in the UCD matches the change count in the UCD of the new channel. The CMTS **MUST** ensure that the Upstream Channel ID for the new channel is different than the Channel ID for the old channel. The Ranging Required parameter in the new UCD does not apply in this context, since the functionality is covered by the Initialization Technique TLV.

If the length of the Type 46 TLV exceeds 254 octets after adding the UCD, more than one Type 46 TLV **MUST** be used to encode the TCC TLV. The UCD may need to be fragmented into two or more Type 46.5 fragments encoded in successive Type 46 TLVs. Each fragment **SHOULD** be the largest possible that fits into the space available in its parent Type 46 TLV. The CM reconstructs the UCD Substitution by concatenating the contents (value of the TLV) of successive Type 46.5 fragments in the order in which they appear in the Type 46 TLV fragment sequence of the TCC. For example, the first byte following the length field of the second Type 46.5 fragments is treated as if it immediately follows the last byte of the first Type 46.5 fragment.

C.1.5.1.6 Ranging SID

When present, this TLV provides a SID value to be used by the CM when performing unicast ranging. The CMTS **MUST** include this TLV if the Upstream Channel Action is Add, Change or Replace.

Type	Length	Value
46.6	2	SID to be used for ranging (lower 14-bits of 16-bit field)

C.1.5.1.7 Initialization Technique

When present, this TLV allows the CMTS to direct the CM as to what level of re-initialization it MUST perform before it can commence communications on the new channel.

The CMTS MAY include the Initialization Technique encoding within a TCC when the Upstream Channel Action is Add, Replace or Re-Range. The CMTS MUST NOT include the Initialization Technique encoding within a TCC when the Upstream Channel Action is No action, Change or Delete. The CM MUST observe the Initialization Technique encoding if it is specified within a TCC when the Upstream Channel Action is Add, Replace or Re-Range.

Type	Length	Value
46.7	1	1 = Perform broadcast initial ranging on new channel before normal operation 2 = Perform unicast ranging on new channel before normal operation 3 = Perform either broadcast or unicast ranging on new channel before normal operation 4 = Use new channel directly without reinitializing or ranging 0, 5 to 255: reserved

If this TLV is not present and ranging is required on a channel, the CM MUST perform broadcast initial ranging on the channel before normal operation.

C.1.5.1.8 Ranging Parameters

The CMTS MAY include the Ranging Parameters TLV within the TCC when the Upstream Channel Action is Add or Replace. The CMTS MUST include the Ranging Parameters TLV within the TCC when the Upstream Channel Action is Add or Replace and the Initialization Technique has a value of "2", "3" or "4". The CMTS MUST NOT include the Ranging Parameters encoding within a TCC when the Upstream Channel Action is No action, Change, Delete or Re-range.

The CM MUST observe this TLV. The value field of this TLV contains a series of sub-types describing parameters to be used when initializing on the channel being added or replaced.

Type	Length	Value
46.8	N	

C.1.5.1.8.1 Ranging Reference Channel ID

This TLV MUST be included exactly once in the Ranging Parameters TLV. It provides the ID of a channel whose timing and power values are used as the references for the corresponding offsets.

If the Initialization Technique has a value of 2, 3 or 4, the CM MUST use the result of the accumulated frequency adjustments made on the channel designated as the Reference Channel to calculate proportional frequency offsets for any channels being added or replaced by the TCC.

For example if the Reference Channel UCD frequency is 10 MHz and the CM has been given ranging adjustments increasing the frequency by 100 Hz for that channel, the CM would use a scale factor of 100/10E6 for setting the Transmit Frequency Offset for the channels to be added. Continuing the example, if a channel is being added with a UCD frequency of 20 MHz, the CM would set the Initial Tx Frequency to 20 MHz + (20 MHz * 100/10E6) = 20,0002 MHz rounded to Hz for transmitting on the new channel. The CM is not expected to set its Tx Frequency to fractional Hz.

Subtype	Length	Value
46.8.1	1	1 to 255: Upstream Channel ID

C.1.5.1.8.2 Timing Offset, Integer Part

When present, this TLV provides the value, as an offset from the reference channel, to set the Ranging Offset of the burst transmission for the new channel so that bursts arrive at the expected mini-slots time at the CMTS. The units are $(1/10,24 \text{ MHz}) = 97,65625 \text{ ns}$. A negative value implies the Ranging Offset is to be decreased, resulting in later times of transmission at the CM. The CMTS MUST include this TLV within the TCC when the Upstream Channel Action is Add or Replace and the Initialization Technique has a value of "2", "3" or "4".

The CMTS does not include the timing offset necessary to compensate for differences in modulation rate (Timing Offset for Modulation Rate Changes table in [12]) between the Ranging Reference Channel and the upstream channels being added or replaced in the value of this TLV. The CMTS does not include the timing offset necessary to compensate for differences in the pre-equalizer main tap location between the Ranging Reference Channel and the upstream channels being added or replaced in the value of this TLV. The CM MUST apply this TLV in addition to the timing offset necessary to compensate for differences in modulation rate and pre-equalizer main tap location between the Ranging Reference Channel and the upstream channels being added or replaced.

Subtype	Length	Value
46.8.2	4	TX timing offset adjustment (signed 32-bit, units of (6,25 micros/64))

C.1.5.1.8.3 Timing Offset, Fractional Part

When present, this TLV provides a higher resolution timing adjust offset to be appended to Timing Adjust, Integer Part for the new channel, compared to the reference channel. The units are $(1/(256*10,24 \text{ MHz})) = 0,3814697265625 \text{ ns}$. This parameter provides finer granularity timing offset information for transmission in S-CDMA mode.

Subtype	Length	Value
46.8.3	1	TX timing fine offset adjustment. 8-bit unsigned value specifying the fine timing adjustment in units of $1/(256*10,24 \text{ MHz})$.

C.1.5.1.8.4 Power Offset

When present, this TLV provides the transmission power level, as an offset from the reference channel that the CM is to use on the new channel in order that transmissions arrive at the CMTS at the desired power. The CMTS MUST include this TLV within the TCC when the Upstream Channel Action is Add or Replace and the Initialization Technique has a value of "2", "3" or "4".

Subtype	Length	Value
46.8.4	1	TX power offset adjustment (signed 8-bit, 1/4-dB units)

C.1.5.1.8.5 Frequency Offset

This TLV is deprecated. The CMTS MUST NOT include this TLV in the TCC encodings. The CM MUST ignore this TLV.

Subtype	Length	Value
46.8.5	2	Deprecated - formerly the TX frequency offset adjustment (signed 16-bit Hz units)

C.1.5.1.9 Dynamic Range Window

When present, this TLV specifies the value for the top of the Dynamic Range Window ($P_{\text{load_min_set}}$) [12]. The CMTS MUST include this TLV in the REG-RSP-MP if Multiple Transmit Channel Mode is to be enabled.

Because it is not associated with a single upstream channel, the Dynamic Range Window TLV can only be included when the Upstream Channel ID is "0". When the value of the Dynamic Range Window is changing, the Upstream Channel Action is "change". When the value of the Dynamic Range Window is included in the TCC Encodings but not changing, the Upstream Channel Action is "no action" (for example, the CMTS includes the Dynamic Range Window value in a RNG-RSP prior to registration and includes the same Dynamic Range Window in the REG-RSP-MP).

Subtype	Length	Value
46.9	1	Dynamic Range Window - Pload_min_set expressed in units of 1/4 dB below P _{hi} [12]

NOTE: During normal operation the CMTS controls the CM's Dynamic Range Window value using the RNG-RSP message. Prior to registration, the CM does not need a Dynamic Range Window value. The CM requires a value for the Dynamic Range Window when operating in Multiple Transmit Channel Mode.

C.1.5.1.10 TCC Error Encodings

This TLV is included to report the status of or any errors with, the action directed in the TCC.

Subtype	Length	Value
46.254	n	

C.1.5.1.10.1 Reported Parameter

The value of this parameter identifies the subtype of a TCC that is being reported. A TCC Error Set MUST have exactly one Reported Parameter TLV within a given TCC Error Encoding.

Subtype	Length	Value
46.254.1	n	TCC subtype

If the length is one, then the value is the single-level subtype (for example, a value of 0x06 indicates the Ranging SID (clause C.1.5.1.6)). If the length is two, then the value is the multi-level subtype, with the first byte representing the TCC subtype and the second byte representing the next level subtype (for example, a value of 0x0804 indicates the Power Offset within the Ranging Parameters (clause C.1.5.1.8.4)).

C.1.5.1.10.2 Error Code

This parameter indicates the status of the operation. A non-zero value corresponds to the Confirmation Code as described in clause C.4. A TCC Error Set MUST have exactly one Status Code within a given TCC Status Encoding.

Subtype	Length	Value
46.254.2	1	Confirmation Code

C.1.5.1.10.3 Error Message

This subtype is optional in the TCC Error Set. If present, it indicates a text string to be displayed on the CMTS console and/or log that further describes a rejected TCC operation. A TCC Error Set MAY have zero or one Error Message subtypes within a given TCC Error Encoding.

Subtype	Length	Value
46.254.3	n	Zero-terminated string of ASCII characters

C.1.5.2 Service Flow SID Cluster Assignments

This TLV contains an SFID and channel-to-SID mappings within SID Clusters to be used by the service flow. When present, this TLV MUST be included by the CMTS exactly once per Service Flow.

This TLV can be used in Registration, Dynamic Service Add and Dynamic Bonding Change MAC Management Messages. The CMTS MUST NOT include this TLV in Dynamic Service Change MAC Management Messages.

Type	Length	Value
47	N	Service Flow SID Cluster Assignments

C.1.5.2.1 SFID

The SFID associated with the SID Cluster. This TLV MUST be included exactly once in a Service Flow SID Cluster Assignment.

Type	Length	Value
47.1	4	Service Flow ID

C.1.5.2.2 SID Cluster Encoding

This TLV contains a service flow identifier, the channel-to-SID mappings of the SID clusters associated to the service flow and the service flow's SID Cluster switchover criteria. When present, this TLV MUST be included by the CMTS exactly once for each SID Cluster assigned to the service flow.

Type	Length	Value
at	N	SID Cluster Encodings

C.1.5.2.2.1 SID Cluster ID

This TLV contains the SID Cluster ID in the range of 0 to 7. The CMTS MUST include this encoding exactly once per SID Cluster encoding. The CMTS MUST assign values in the range of 0 to M-1 where M is the number of SID Clusters per Service Flow supported by the CM.

Subtype	Length	Value
47.2.1	1	SID Cluster ID

C.1.5.2.2.2 SID-to-Channel Mapping

When present, this TLV MUST be included by the CMTS once per channel. This TLV contains the mapping of a channel ID to SID in the SID Cluster. The value field consists of three sub-TLVs. When this TLV is present, the CMTS MUST include each sub-TLV exactly once.

Subtype	Length	Value
47.2.2	10	Sub-TLVs as described below

C.1.5.2.2.2.1 SID-to-Channel Mapping: Upstream Channel ID

This subtype indicates the channel ID on which a SID is being mapped.

Subtype	Length	Value
47.2.2.1	1	Upstream Channel ID

C.1.5.2.2.2.2 SID-to-Channel Mapping: SID

This subtype gives the SID which is being mapped to the channel indicated in subtype 47.2.2.1.

Subtype	Length	Value
47.2.2.2	2	2-byte SID (lower 14 bits of 16-bit field)

C.1.5.2.2.3 SID-to-Channel Mapping: Action

This subtype indicates whether the SID indicated in subtype 47.2.2.2 is being added or deleted.

Subtype	Length	Value
47.2.2.3	1	Action:
		1 = add
		2 = delete
		0, 3-255 = reserved

C.1.5.2.3 SID Cluster Switchover Criteria

This TLV contains the SID Cluster Switchover criteria for use by the service flow. The CMTS MAY include this sub-TLV. If the CMTS includes this sub-TLV, it MUST NOT repeat it more than once for a service flow. If the CMTS includes this sub-TLV, it MUST define within it at least one SID Cluster switchover criteria.

Type	Length	Value
47.3	N	SID Cluster Switchover Criteria

C.1.5.2.3.1 Maximum Requests per SID Cluster

This is the maximum number of requests that a CM can make with a given SID Cluster before it must switch to a different SID Cluster to make further requests. The CMTS MAY include this sub-TLV. The CMTS MUST NOT include this TLV more than once within a SID Cluster Switchover Criteria sub-TLV.

Type	Length	Value
47.3.1	1	1 to 255 requests
		0 = unlimited

A value of 0 represents no limit. If not present, a default value of 0 is used.

C.1.5.2.3.2 Maximum Outstanding Bytes per SID Cluster

This is the maximum number of bytes for which a CM can have requests outstanding on a given SID Cluster. If this many bytes are outstanding and further requests are required, the CM must switch to a different SID Cluster if one is available. If a different SID Cluster is not available, then the CM will stop requesting until there are no bytes outstanding for which the acknowledgement time has not passed. The CMTS MAY include this sub-TLV. The CMTS MUST NOT include this TLV more than once within a SID Cluster Switchover Criteria sub-TLV.

Type	Length	Value
47.3.2	4	1 to 4 294 967 295 bytes
		0 = unlimited

A value of 0 represents no limit. If not present, a default value of 0 is used.

C.1.5.2.3.3 Maximum Total Bytes Requested per SID Cluster

This is the maximum total number of bytes a CM can have requested using a given SID Cluster before it must switch to a different SID Cluster to make further requests. The CMTS MAY include this sub-TLV. The CMTS MUST NOT include this TLV more than once within a SID Cluster Switchover Criteria sub-TLV.

Type	Length	Value
47.3.3	4	1 to 4 294 967 295 bytes
		0 = unlimited

A value of 0 represents no limit. If not present, a default value of 0 is used.

C.1.5.2.3.4 Maximum Time in the SID Cluster

This is the maximum time in milliseconds that a CM may use a particular SID Cluster before it must switch to a different SID Cluster to make further requests. The CMTS MAY include this sub-TLV. The CMTS MUST NOT include this TLV more than once within a SID Cluster Switchover Criteria sub-TLV.

Type	Length	Value
47.3.4	2	1 to 65 535 milliseconds
		0 = unlimited

A value of 0 represents no limit. If not present, a default value of 0 is used.

C.1.5.3 CM Receive Channel (RCP/RCC) Encodings

The CM includes one or more Receive Channel Profile (RCP) Encodings in its Registration Request to describe the physical layer components that permit it to receive multiple downstream channels. The CMTS returns to the CM in a Registration Response a Receive Channel Configuration (RCC) Encoding that configures the physical layer components to certain frequencies and, if necessary, to certain interconnections between those components.

After a CM has registered, the CMTS changes the set of downstream channels received by a CM with a Dynamic Bonding Change Request (DBC-REQ) message that contains a Receive Channel Configuration Encoding.

The Receive Channel Profile Encoding and Receive Channel Configuration Encoding contain many sub-types in common. In this annex, a Receive Channel Profile subtype is denoted as "48.x" and a Receive Channel Configuration subtype is denoted as "49.x".

Type	Length	Value
48	N	Receive Channel Profile Subtype TLVs
49	N	Receive Channel Configuration Subtype TLVs

The CM MUST support these TLVs. The CM MAY repeat the RCP TLV in a Registration-Request to describe multiple Receive Channel Profiles. The CMTS MUST support these TLVs. The CMTS MUST silently ignore invalid RCP encodings. The CMTS MUST silently ignore unknown RCP subtype encodings and process known RCP subtype encodings normally.

The CM MUST fragment an RCP encoding that exceeds 255 bytes in length (note that as per clause 6.4.28.1.4, the CM only sends these fragmented RCPs if the CMTS indicates that it can support them). The CMTS MUST support reception of fragmented RCPs.

The CMTS MUST support the ability to fragment RCCs that are greater than 255 bytes in length. The CMTS MUST NOT fragment an RCC that is 255 bytes or less in length. The CM MUST support the reception of a fragmented RCC from the CMTS.

The CMTS MUST NOT transmit a fragmented RCC to a CM that advertises a Multiple Receive Channel Support capability of less than 8 (clause C.1.3.1.29).

If an RCP is fragmented, the CM MUST fragment the RCP at sub-TLV boundaries within the Receive Channel Profile TLV, TLV 48. This means that an RCP fragment contains complete RCP sub-TLVs.

If an RCC is fragmented, the CMTS MUST fragment the RCC at sub-TLV boundaries within the Receive Channel configuration TLV, TLV 49. This means that an RCC fragment contains complete RCC sub-TLVs.

C.1.5.3.1 RCP-ID

In an RCP, the RCP-ID identifies the RCP being described. A REG-REQ-MP may have multiple RCP Encodings that describe different logical profiles for configuring the physical interface of the CM.

A Receive Channel Configuration has a single RCP-ID that assigns the CM to use a particular Receive Channel Profile that it supports. The CMTS MAY change the assigned RCP-ID for a CM in a DBC-REQ to the CM. The CM MUST support a change of RCP-ID communicated in a DBC-REQ message.

The CM MUST include the RCP-ID sub-TLV as the first sub-TLV and exactly once within each instance of TLV 48 (RCP) that it transmits. In other words, each RCP fragment will start with the RCP-ID sub-TLV and unfragmented RCPs will also start with the RCP-ID.

The CMTS MUST include the RCP-ID sub-TLV as the first sub-TLV and exactly once within each instance of TLV 49 (RCC) that it transmits. In other words, each RCC fragment will start with the RCP-ID sub-TLV and unfragmented RCCs will also start with the RCP-ID.

Type	Length	Value
48.1	5	Bytes 0, 1, 2: Organization Unique ID Bytes 3, 4: OUI-specific profile ID
49.1	5	Assigned RCP-ID

C.1.5.3.2 RCP Name

This parameter defines a human-readable, descriptive name for the Receive Channel Profile. The RCP Name is assigned by the vendor and is not guaranteed to be globally unique. It is recommended that the vendor assign RCP Names uniquely within an OUI. The CM MAY include the RCP Name encoding in an RCP encoding.

Type	Length	Value
48.2	1..15	Informational DisplayString corresponding to RCP-ID

C.1.5.3.3 RCP Center Frequency Spacing

This parameter defines the interval between center frequencies in a Receive Module. The CM MUST include the RCP Center Frequency Spacing TLV in a verbose RCP encoding. The CM MUST NOT include the RCP Center Frequency Spacing TLV in a non-verbose RCP encoding.

Type	Length	Value
48.3	1	6 = 6 MHz channels 8 = 8 MHz channels

C.1.5.3.4 Receive Module Encoding

This TLV describes a Receive Module of the CM. A Receive Module is often configured to be a block of adjacent center channel frequencies at the center frequency spacing of the RCP.

Each Receive Module Encoding consists of multiple subtypes.

The CM MAY include the Receive Module Encoding TLV in a verbose RCP encoding. The CM MUST NOT include the Receive Module Encoding TLV in a non-verbose RCP encoding. In the RCC, the CMTS MUST include all Receive Module encodings associated with the Receive Channels configured in the RCC.

Type	Length	Value
48.4	N	Receive Module Capability
49.4	N	Receive Module Assignment

C.1.5.3.4.1 Receive Module Index

This is signaled by the CM in an RCP and the CMTS in an RCC to identify a Receive Module. This parameter is required to be present exactly once in each Receive Module Encoding. The CM MUST include exactly one Receive Module Index in a Receive Module Encoding of an RCP Encoding. The CMTS MUST include exactly one Receive Module Index in a Receive Module Encoding of an RCC Encoding.

Type	Length	Value
48.4.1	1	Receive Module index being described, starting from 1
49.4.1	1	Receive Module index being assigned

C.1.5.3.4.2 Receive Module Adjacent Channels

If the Receive Module corresponds to a block of adjacent channel center frequencies, this parameter provides the number of such channels in the block. The CM MAY include the Receive Module Adjacent Channels TLV in a Receive Module encoding of an RCP encoding.

Type	Length	Value
48.4.2	1	Number of adjacent center frequencies in the Receive Module channel block

C.1.5.3.4.3 Receive Module Channel Block Range

DOCSIS defines various downstream frequency ranges over which a CM may be capable of operating. This parameter is used to indicate the frequency range over which a Receive Module can be tuned. Further, the CM may not be able to tune a channel block anywhere in its full downstream frequency range. In either case, this parameter indicates the limited range of the channel block in terms of a minimum of the first center frequency of the channel block and a maximum of the last center frequency of the channel block. This parameter is encoded with two required subtypes.

The CM MAY include the Receive Module Channel Block Range TLV in a Receive Module Encoding of an RCP Encoding. For RCPs indicating a RCP Center Frequency Spacing of 6 MHz, the absence of this TLV is equivalent to a Receive Module Minimum Center Frequency of 111 MHz and a Receive Module Maximum Center Frequency of 867 MHz. For RCPs indicating a RCP Center Frequency Spacing of 8 MHz, the absence of this TLV is equivalent to a Receive Module Minimum Center Frequency of 112 MHz and a Receive Module Maximum Center Frequency of 858 MHz.

Type	Length	Value
48.4.3	12	The Minimum Center Frequency and Maximum Center Frequency subtypes as described immediately below.

C.1.5.3.4.3.1 Receive Module Minimum Center Frequency

Type	Length	Value
48.4.3.1	4	Minimum center frequency (Hz) of the first channel of the block

C.1.5.3.4.3.2 Receive Module Maximum Center Frequency

Type	Length	Value
48.4.3.2	4	Maximum center frequency (Hz) of the last channel of the block

C.1.5.3.4.4 Receive Module First Channel Center Frequency Assignment

This subtype is included only in a Receive Channel Configuration (RCC) to assign a Receive Module corresponding to a block of adjacent center frequencies to a particular point in the spectrum. When the Receive Module Adjacent Channels TLV is present in a Receive Module associated with an assigned Receive Channel, the CMTS MUST include a Receive Module First Channel Center Frequency Assignment TLV in its RCC to the CM. The CMTS MUST NOT assign a First Channel Center Frequency such that any center frequency in the channel block falls outside the frequency range limits communicated in the Receive Module Channel Block Range. The CMTS MUST assign the First Channel Center Frequency to be a multiple of 62 500 Hz.

Type	Length	Value
49.4.4	4	Assigned center frequency of the first channel of the Receive Module channel block, in Hz.

C.1.5.3.4.5 Receive Module Resequencing Channel Subset Capability

This parameter, if present in a Receive Module Encoding, signals that the Receive Module represents a subset of Receive Channels of the CM within which resequencing can be performed. If omitted, the CMTS assumes that any subset of Receive Channels of the CM may be signaled as a Resequencing Channel List for a DSID. The CM MAY include one or more Resequencing Channel Subset encodings in a Receive Module encoding of a RCP. The CM MUST NOT signal more than one Resequencing Channel Subset encoding for any Receive Channel.

Type	Length	Value
48.4.5	N	BITS Encoding with bit position N set to 1 if Receive Channel N is a part of the subset within which resequencing can be performed. Bit position 0 (the most significant bit) is unused and must be zero.

C.1.5.3.4.6 Receive Module Connectivity

This parameter, if present in an RCP, indicates via a bit map the set of other "higher-layer" Receive Modules to which the currently described Receive Module may attach. If more than one higher-layer Receive Module is signaled, the CMTS MUST select only one of them and include a Receive Module Connectivity subtype in an RCC that indicates the single other higher-layer Receive Module that it selected. The CM MAY include the Receive Module Connectivity TLV in a Receive Module encoding of a RCP.

Type	Length	Value
48.4.6	N	BITS Encoding with bit K set to 1 for each Receive Module Index K to which the currently described Receive Module may connect. Bit 0 is the most significant bit.
49.4.6	N	BITS Encoding with one bit set for the Receive Module to which the current Receive Module is assigned to attach. Bit 0 is the most significant bit.

C.1.5.3.4.7 Receive Module Common Physical Layer Parameter

This parameter, if present in an RCP, indicates which physical layer parameters must be the same for all Receive Channels connected to the Receive Module. The CM MAY include the Receive Module Common Physical Layer Parameter TLV in a Receive Module encoding of a RCP.

Type	Length	Value
48.4.7	N	BITS Encoding indicating what parameters must be the same: Bit Position 0 (0x80): QAM Modulation Order Bit Position 1 (0x40): Interleave

C.1.5.3.5 Receive Channels

Receive Channels (RCs) represent individual demodulators. Receive Channels may be associated with a single position within a Receive Module's channel block.

The CM MUST include at least one Receive Channel subtype in each Receive Channel Profile Encoding. The CMTS MUST assign at least one Receive Channel subtype in each Receive Channel Configuration Encoding. The CMTS is not required to send a Receive Channel subtype in the Receive Channel Configuration for every Receive Channel subtype present in the Receive Channel Profile.

Type	Length	Value
48.5	N	Receive Channel (RC) capable of being assigned
49.5	N	Receive Channel assigned by CMTS

C.1.5.3.5.1 Receive Channel Index

The CM MUST include exactly one Receive Channel Index in each Receive Channel Encoding in an RCP encoding. The CMTS MUST include exactly one Receive Channel Index in each Receive Channel Encoding in an RCC encoding.

Type	Length	Value
48.5.1	1	RC Index within the RCP
49.5.1	1	RC Index within the RCC

C.1.5.3.5.2 Receive Channel Connectivity

This parameter, if present in an RCP, indicates via a bit map the non-null set of Receive Modules to which the Receive Channel may attach. If the Receive Channel is not connected to any Receive Module, the CM MUST omit this parameter. When present in an RCP, the CMTS MUST select a single Receive Module and include a Receive Channel Connectivity subtype in an RCC that indicates the single Receive Module that it selected. For RCPs indicating a RCP Center Frequency Spacing of 6 MHz, the absence of this TLV indicates that the Receive Channel can be assigned to any center frequency between 111 MHz and 867 MHz. For RCPs indicating a RCP Center Frequency Spacing of 8 MHz, the absence of this TLV indicates that the Receive Channel can be assigned to any center frequency between 112 MHz and 858 MHz.

Type	Length	Value
48.5.2	N	Receive Channel Connectivity Capability. BITS encoding with bit position K set to 1 when RC can connect to Receive Module Index K. Bit position 0 is the most significant bit.
49.5.2	N	Receive Channel Connectivity Assignment. BITS encoding with only 1 bit position K set indicating the assigned connection of the RC to the Receive Module with index K. Bit position 0 is the most significant bit.

C.1.5.3.5.3 Receive Channel Connected Offset

When an RCP Receive Channel Connectivity indicates that the RC is connected to a single Receive Module corresponding to a block of channels, this parameter can be used to indicate a fixed position that this Receive Channel occupies in that Receive Module. The position of 1 indicates the first (i.e. lowest frequency) channel in the Receive Module. The CM MAY include the Receive Channel Connected Offset in a Receive Channel encoding of an RCP encoding.

Type	Length	Value
48.5.3	1	Assigned (1-based) position with the channel block of a single Receive Module

C.1.5.3.5.4 Receive Channel Center Frequency Assignment

The CMTS MUST include the Receive Channel Center Frequency Assignment TLV in a Receive Channel encoding of an RCC encoding to assign a particular center frequency to a Receive Channel. The CMTS MUST assign the Center Frequency as a multiple of 62 500 Hz.

Type	Length	Value
49.5.4	4	Assigned center frequency of the channel, in Hz.

C.1.5.3.5.5 Receive Channel Primary Downstream Channel Indicator

This subtype is included in a Receive Channel Profile (RCP) or Receive Channel Configuration (RCC) to control assignment of the CM's Primary Downstream Channel.

Type	Length	Value
48.5.5	1	A value of 1 indicates that the Receive Channel is capable of operating as the CM's primary downstream channel. A value of 0 indicates that the Receive Channel is not capable of operating as the CM's primary downstream channel. The CM MUST signal at least one Receive Channel as being capable of operating as the primary downstream channel. If omitted, the default is 0.
49.5.5	1	A value of 1 indicates that the channel is assigned to be the CM's primary downstream channel. A value of 0 indicates that the channel is not assigned to be the CM's primary downstream channel. The CMTS MUST assign a single Receive Channel as the CM's primary downstream channel. If omitted, the default is 0.

C.1.5.3.6 Partial Service Downstream Channels

This subtype is used to provide the CMTS a list of the downstream channels that could not be acquired by the CM as a result of a REG-RSP-MP or a DBC-REQ. The CM MUST include the Partial Service Downstream Channels TLV if there were no errors in the RCC, but it was unable to acquire all of the downstream channels it was directed to by the RCC.

Type	Length	Value
49.6	N	List of N 1-byte downstream channel IDs that could not be acquired.

C.1.5.3.7 Receive Channel Profile/Configuration Vendor Specific Parameters

The CM MAY include Vendor Specific Parameters in a manufacturer-specific RCP encoding. The CMTS MAY include Vendor Specific Parameters in an RCC encoding assigned to a manufacturer-specific profile.

A valid Vendor Specific Parameter Encoding is encoded as a set of subtypes with the first subtype providing the Vendor Identifier subtype (see clause C.1.3.2).

Type	Length	Value
48.43	N	Vendor Specific Parameters
49.43	N	Vendor Specific Parameters

C.1.5.3.8 RCC Error Encodings

This TLV is included to report the status of or any errors with, the actions directed in the RCC.

Type	Length	Value
49.254	n	

C.1.5.3.8.1 Receive Module or Receive Channel

The value of this parameter identifies whether the error being reported applies to a Receive Module or a Receive Channel. An RCC Error Set MUST have exactly one Receive Module or Receive Channel TLV within a given RCC Error Encoding.

Type	Length	Value
49.254.1	1	4 = Receive Module 5 = Receive Channel 0 to 3, 6 to 255 = Reserved

C.1.5.3.8.2 Receive Module Index or Receive Channel Index

The value of this parameter identifies the Receive Module Index or Receive Channel Index that is being reported. An RCC Error Set MUST have exactly one Receive Module Index or Receive Channel Index TLV within a given RCC Error Encoding.

Subtype	Length	Value
49.254.2	1	Receive Module Index or Receive Channel Index

C.1.5.3.8.3 Reported Parameter

The value of this parameter identifies the subtype of a Receive Module or Receive Channel that is being reported. An RCC Error Set MUST have exactly one Reported Parameter TLV within a given RCC Error Encoding.

Subtype	Length	Value
49.254.3	1	Receive Module or Receive Channel Subtype

C.1.5.3.8.4 Error Code

This parameter indicates the status of the operation. A non-zero value corresponds to the Confirmation Code as described in clause C.4. An RCC Error Set MUST have exactly one Error Code within a given RCC Error Encoding.

Subtype	Length	Value
49.254.4	1	Confirmation Code

C.1.5.3.8.5 Error Message

This subtype is optional in the RCC Error Set. If present, it indicates a text string to be displayed on the CMTS console and/or log that further describes a rejected RCC operation. An RCC Error Set MAY have zero or one Error Message subtypes within a given RCC Error Encoding.

Subtype	Length	Value
49.254.5	n	Zero-terminated string of ASCII characters

C.1.5.4 DSID Encodings

The value of this field is used by the CMTS to provide the CM with the DSID encodings assigned by the CMTS. It can be used in Registration and DBC MAC Management Messages.

Type	Length	Value
50	N	DSID Encodings

The CMTS MAY include multiple instances of these TLVs.

C.1.5.4.1 Downstream Service Identifier (DSID)

The value of this field is used by the CMTS to provide the CM with the DSID assigned by the CMTS.

Type	Length	Value
50.1	3	DSID (1 to 1 048 575)

The CMTS MUST include this TLV.

C.1.5.4.2 Downstream Service Identifier Action

The value of this field is used by the CMTS to inform the CM as to whether it is adding, changing or deleting the DSID.

Type	Length	Value
50.2	1	0 = Add 1 = Change 2 = Delete 3 to 255: Reserved

The CMTS MUST include this sub-TLV with any DSID encoding.

C.1.5.4.3 Downstream Resequencing Encodings

The value of this field specifies the downstream resequencing encodings assigned by the CMTS.

Type	Length	Value
50.3	N	Encoded resequencing attributes

The CMTS MUST include this TLV if adding or changing a resequencing DSID. The CMTS MUST NOT include this TLV if the DSID is a not a resequencing DSID.

C.1.5.4.3.1 Resequencing DSID

The value of this field is used by the CMTS to notify the CM that the DSID is being used for resequencing.

Subtype	Length	Value
50.3.1	1	1 = DSID is a resequencing DSID 0, 2 to 255: Reserved

The CMTS MUST include this sub-TLV.

C.1.5.4.3.2 Downstream Resequencing Channel List

The value of the field is used by the CMTS to provide the CM with a list of downstream channels associated with the DSID for reassembly.

Subtype	Length	Value
50.3.2	n	DCID [1]. DCID [2], ... , DCID[n]

The CMTS MAY include this sub-TLV. If rapid loss detection is desired for a subset of channels within the Receive Channel Set, the CMTS MUST include this sub-TLV. If this sub-TLV is present, the CM MUST perform rapid loss detection on the set of downstream channels indicated by this sub-TLV. If this sub-TLV is not present, the CM MUST associate all of the channels in the Receive Channel Set with the DSID for rapid loss detection.

C.1.5.4.3.3 DSID Resequencing Wait Time

The value of the field is used by the CMTS to provide the CM with the value of the DSID Resequencing Wait Time in units of 100 μ s.

Subtype	Length	Value
50.3.3	1	1 to 180

The CMTS MAY include this sub-TLV. If this TLV is not included for a resequencing DSID, the CM MUST assume the maximum DSID Resequencing Wait Time value defined in annex B.

C.1.5.4.3.4 Resequencing Warning Threshold

The usage of this field is described in clause 8.2.3.

Subtype	Length	Value
50.3.4	1	0 to 179

The CMTS MAY include this sub-TLV. If included, the value of Resequencing Warning Threshold MUST be less than the value of DSID Resequencing Wait Time. If this TLV is not included for a resequencing DSID or is included with the value 0, the CM MUST assume that threshold counting and reporting is disabled.

C.1.5.4.3.5 CM-STATUS Maximum Event Hold-Off Timer for Sequence Out-of-Range Events

The value of this field is used by the CMTS to provide the CM with the value of the hold-off timer for the out-of-range events in units of 20 ms.

Subtype	Length	Value
50.3.5	2	CM-STATUS Hold-off Timer for Out-of-Range Events (in 20 ms)

The CMTS MAY include this sub-TLV. If this TLV is not included for a resequencing DSID, the CM MUST use the STATUS Backoff Timer value communicated to the CM in the MDD message.

C.1.5.4.4 Multicast Encodings

The value of this field specifies the multicast encodings assigned by the CMTS to a DSID.

Type	Length	Value
50.4	N	Encoded multicast attributes

C.1.5.4.4.1 Client MAC Address Encodings

The value of this field is used by the CMTS to provide the CM with the client MAC address(es) joining or leaving the multicast group.

Subtype	Length	Value
50.4.1	N	Client MAC address encodings

The CMTS MAY include multiple instances of this sub-TLV. The CMTS MUST include exactly one of the client MAC address action and client MAC address TLV encodings for each instance of this TLV. See clause 11.5.1.2.2 for the interaction with the Multicast CMIM.

C.1.5.4.4.1.1 Client MAC Address Action

The value of this field is used by the CMTS to inform the CM as to whether it is to add or delete the client MAC address.

Subtype	Length	Value
50.4.1.1	1	0 = Add 1 = Delete 2 to 255: Reserved

C.1.5.4.4.1.2 Client MAC Address

The value of this field is used by the CMTS to provide the CM with the source MAC address joining or leaving the multicast group associated with the group flow label.

Subtype	Length	Value
50.4.1.2	6	Client MAC Address

C.1.5.4.4.2 Multicast CM Interface Mask

This field is used by the CMTS to provide a bit mask representing the interfaces of the CM to which the CM is to forward multicast traffic associated with the DSID. Each bit of CM interface mask corresponds to an interface, logical or physical. By convention, bit position 0 corresponds to the CM's IP stack, even though it is not an actual interface.

For example, a Multicast CMIM intended to match all of the external CPE interfaces of a CM has a CMIM value setting bits 1 and 5 to 15, i.e. an encoding of either 0x47FF or 0x47FF0000. Either value is valid.

Subtype	Length	Value
50.4.2	N	BITS Encoded bit map with bit position K representing eCM logical interface index value K. Bit position 0 represents the eCM "self" host itself. Bit position 0 is the most significant bit of the most significant octet. The Embedded DOCSIS specification [eDOCSIS] defines the interface index assignments. For information purposes, current assignments include: Bit 0 (0x80): CM's IP stack Bit 1 (0x40): primary CPE Interface (also ePS or eRouter) Bit 2 (0x20) RF interface Bits 3,4 reserved Bits 5..15 (0x07 FF) Other CPE Interfaces Bits 16-31, Logical CPE Interfaces for eSAFE hosts. Current assignments include: Bit 16 (0x00 00 80) PacketCable-EMTA Bit 17 (0x00 00 40) eSTB-IP Bit 18 (0x00 00 20) reserved Bits 19..31 (0x00 00 1F FF) Other eSAFE interfaces

The CMTS MAY include exactly one instance of this sub-TLV. See clause 11.5.1.2.2 for the interaction with the Client MAC Address Encodings.

C.1.5.4.4.3 Multicast Group MAC Addresses Encodings

The value of this field is used by the CMTS to provide the CM with the multicast group MAC address(es) (GMACs) of the multicast group. In most cases, the CMTS will provide one GMAC.

Type	Length	Value
50.4.3	N	GMAC [1], GMAC [2], ... , GMAC[n]

If the CMTS has confirmed support for GMAC explicit multicast DSID filtering in the modem capabilities, the CMTS MUST include this sub-TLV. If the CMTS has confirmed support for GMAC promiscuous multicast DSID filtering in the modem capabilities, the CMTS MUST NOT include this sub-TLV.

C.1.5.4.4.4 Payload Header Suppression Encodings

The value of this field is used by the CMTS to provide the CM with the parameters associated with Payload Header Suppression. A valid Multicast Encoding contains no more than one Payload Header Suppression encoding (see clause C.2.2.8).

Subtype	Length	Value
50.4.26.x	N	PHS Encodings (clause C.2.2.8)

If the CMTS includes Payload Header Suppression Encodings in the DSID Multicast Encodings of a REG-RSP, REG-RSP-MP or DBC-REQ, the CMTS MUST include the DBC Action. If the CMTS includes Payload Header Suppression Encodings in the DSID Multicast Encodings of a REG-RSP, REG-RSP-MP or DBC-REQ, the CMTS MUST include the Payload Header Suppression Field (PHSF) and the Payload Header Suppression Size (PHSS) in the Payload Header Suppression Rule Encodings.

If the CMTS includes Payload Header Suppression Encodings in the DSID Multicast Encodings of a REG-RSP, REG-RSP-MP or DBC-REQ, the CMTS MAY include the Payload Header Suppression Mask (PHSM), Payload Header Suppression Verify (PHSV) and the Vendor Specific PHS Parameters in the Payload Header Suppression Rule Encodings.

If the CMTS includes Payload Header Suppression Encodings in the DSID Multicast Encodings of a REG-RSP, REG-RSP-MP or DBC-REQ, the CMTS MUST NOT include the Classifier Reference, Classifier Identifier, Service Flow Reference, Service Flow Identifier or Dynamic Service Change Action. If the CMTS includes Payload Header Suppression Encodings in the DSID Multicast Encodings of a REG-RSP, REG-RSP-MP or DBC-REQ, the CMTS MUST NOT include the Payload Header Suppression Index (PHSI) in the Payload Header Suppression Rule Encodings.

For example, the CMTS uses type 50.4.26.13 encoding to indicate the Dynamic Bonding Change Action and a type 50.4.26.6 encoding to indicate the Payload Header Suppression Rule Encodings.

C.1.5.5 Security Association Encoding

The value of the field is used by the CMTS to provide the CM with a Security Association with which to encrypt downstream traffic. The CMTS MUST transmit valid Security Association Encodings, as described in this clause. A CM MUST reject invalid Security Association Encodings.

A REG-RSP, REG-RSP-MP or DBC-REQ message may contain any number of Security Association Encodings.

Type	Length	Value
51	N	SA Encoding

C.1.5.5.1 SA Action

This field informs the CM as to whether it is to add or delete a Security Association. A valid Security Association Encoding contains exactly one instance of this subtype.

Subtype	Length	Value
51.1	1	0 = Add 1 = Delete 2 to 255: Reserved

C.1.5.5.2 SA-Descriptor

This field provides the SA-Descriptor of the Security Association to be added or deleted. A valid Security Association Encoding contains exactly one instance of this subtype.

This is a compound attribute whose sub-attributes describe the properties of a Security Association. These properties are the SAID, the SA type and the cryptographic suite employed by the Security Association.

The SA-Descriptor and details of its sub-attributes are defined in the DOCSIS 3.0 Security Specifications [15] in the BPKM Attributes clause in the BPKM Protocol chapter. The CMTS MUST implement a 2 byte Length field for the SA-Descriptor (Sub-TLV 51.23). The CM MUST implement a 2 byte Length field for the SA-Descriptor (Sub-TLV 51.23). This differs from the normal MULPI requirement of a 1-byte Length field for a TLV, in order to maintain consistency with the DOCSIS 3.0 Security Specification [15] which defines the Length field for the SA-Descriptor TLV to be 2 bytes long.

Subtype	Length (2 Octets)	Value
51.23	14	SA-Descriptor Sub Attributes

C.1.5.6 Initializing Channel Timeout

This field defines the maximum total time that the CM can spend performing initial ranging on the upstream channels described in the REG-RSP, REG-RSP-MP or DBC-REQ messages. If the CM is still unsuccessful ranging on any channels when this timer expires, it MUST respond with a REG-ACK or DBC-RSP respectively with error messages. The CMTS MUST include this TLV if Broadcast Initial Maintenance is used. If this TLV is not present, the default timeout is used as defined in annex B.

Type	Length	Value
52	2	1 to 65 535 s

C.2 Quality-of-Service-Related Encodings

C.2.1 Packet Classification Encodings

The following type/length/value encodings MUST be used in 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.

NOTE: Unless otherwise stated, the same sub-TLV types are valid for the Upstream Packet Classification Encoding, the Upstream Drop Packet Classification Encoding and the Downstream Packet Classification Encoding top-level TLVs. These type fields are not valid in other encoding contexts.

A classifier MUST contain at least one encoding from clauses C.2.1.6, C.2.1.7, C.2.1.8, C.2.1.9 or C.2.1.10.

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 v4 and IPv6 header fields (clauses C.2.1.6 and C.2.1.10).

C.2.1.1 Upstream Packet Classification Encoding

This field defines the parameters associated with an upstream Classifier.

Type	Length	Value
22	n	

C.2.1.2 Upstream Drop Packet Classification Encoding

This field defines the parameters associated with an Upstream Drop Classifier.

Type	Length	Value
60	n	

C.2.1.3 Downstream Packet Classification Encoding

This field defines the parameters associated with a downstream Classifier.

NOTE: 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.2.1.4 General Packet Classifier Encodings

C.2.1.4.1 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
[22/23/60].1	1	1 to 255

The CM MUST use the Classifier Reference as the Classifier ID when implementing the Upstream Drop Classifiers provided in the configuration file because the CMTS does not provide a Classifier ID in the REG-RSP-MP message.

C.2.1.4.2 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
[22/23/60].2	2	1 to 65 535

C.2.1.4.3 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/REG-REQ-MP) 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 to 65 535

C.2.1.4.4 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/REG-REQ-MP). 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 to 4,294,967,295

C.2.1.4.5 Rule Priority

The value of this field specifies the priority for the Classifier, which is used for determining the classification order. A higher value indicates higher priority.

Classifiers that appear in Configuration files and Registration messages can have priorities in the range 0 to 255. If no Rule Priority is specified in the Registration Request, the CMTS MUST use the default Rule Priority of 0. If no Rule Priority is specified in the Registration Response, the CM MUST use the default Rule Priority of 0. Classifiers that appear in the DSA/DSC message MUST have priorities in the range 64 to 191, with the default value 64.

The Rule Priority of the Upstream QoS Classifier and the Rule Priority of the Upstream Drop Classifier interact. If a packet matches both an Upstream QoS Classifier and an Upstream Drop Classifier, the CM MUST select the Classifier with the higher Rule Priority.

Type	Length	Value
[22/23/60].5	1	

C.2.1.4.6 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.2.1.4.7 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
[22/23/60].7	1	0 - DSC Add Classifier
		1 - DSC Replace Classifier
		2 - DSC Delete Classifier

C.2.1.4.8 CM Interface Mask (CMIM) Encoding

In addition to classifying traffic based on L2/L3/L4 fields in the packet headers, upstream traffic can be classified based on which CM interface received the packet. The CM Interface Mask Encoding provides a bit mask representing the inbound interfaces of the CM for which this classifier applies. Each bit of the CM Interface Mask corresponds to an interface, logical or physical. By convention, bit position 0 corresponds to the CM's IP stack, even though it is not an actual interface.

For example, a CMIM classifier intended to match all of the CPE ports (i.e. external interfaces) of a CM has a CMIM value setting bits 1 and 5-15, i.e. an encoding of either 0x47FF or 0x47FF0000. Either value is valid.

SubType	Length	Value
[22/60].13	N	BITS -Encoded bit map with bit position K representing CM interface index value K. Bit position 0 is the most significant bit of the most significant octet. Refer to [7] for latest logical interface index assignments for eCMs. Bit 0 (0x80): CM's IP stack Bit 1 (0x40): primary CPE Interface (also ePS or eRouter) Bit 2 (0x20) RF interface Bits 3,4 reserved Bits 5..15 (0x07 FF) Other CPE Ports Bits 16-31, embedded logical interfaces. Currently defined interfaces include: Bit 16 (0x00 00 80) PacketCable-eMTA Bit 17 (0x00 00 40) eSTB-IP Bit 18 (0x00 00 20) reserved Bits 19..31 (0x00 00 1F FF) Other eSAFE interfaces

C.2.1.5 Classifier Error Encodings

This field defines the parameters associated with Classifier Errors.

Type	Length	Value
[22/23/60].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, REG-RSP-MP, 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, REG-REQ-MP, DSA-REQ or DSC-REQ message.

On failure, the REG-RSP, REG-RSP-MP, DSA-RSP or DSC-RSP MUST include one Classifier Error Encoding for at least one failed Classifier requested in the REG-REQ, REG-REQ-MP, 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. 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, REG-RSP-MP, 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, REG-REQ-MP, DSA-REQ or DSC-REQ messages.

C.2.1.5.1 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
[22/23/60].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.2.1.5.2 Error Code

This parameter indicates the status of the request. A non-zero value corresponds to the Confirmation Code as described in clause C.4. A Classifier Error Parameter Set MUST have exactly one Error Code within a given Classifier Error Encoding.

Subtype	Length	Value
[22/23/60].8.2	1	Confirmation code

A value of okay (0) indicates that the Classifier request was successful. Since a Classifier Error Parameter Set applies only to errored parameters, this value MUST NOT be used.

C.2.1.5.3 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
[22/23/60].8.3	n	Zero-terminated string of ASCII characters.

NOTE: The length N includes the terminating zero.
Since the entire Classifier Encoding is limited to a total length of 256 bytes (254 bytes + type + length), the maximum length of the error message string is limited by the number of other sub-TLV encodings in the Classifier Encoding.

C.2.1.6 IPv4 Packet Classification Encodings

This field defines the parameters associated with IPv4 packet classification, as well as parameters associated with TCP/UDP packet classification associated with both IPv4 and IPv6. See clause C.2.1.10 for more details.

Type	Length	Value
[22/23/60].9	n	

C.2.1.6.1 IPv4 Type of Service Range and Mask

The values of the field specify the matching parameters for the IPv4 TOS byte range and mask. An IP packet with IPv4 TOS byte value "ip-tos" matches this parameter if $\text{tos-low} \leq (\text{ip-tos AND tos-mask}) \leq \text{tos-high}$. If this field is omitted, then comparison of the IP packet TOS byte for this entry is irrelevant.

Type	Length	Value
[22/23/60].9.1	3	tos-low, tos-high, tos-mask

NOTE: The value 0x3F for tos-mask will exclude the Explicit Congestion Notification [15] bits from the comparison and hence will result in classification based on DSCP [i.5].

C.2.1.6.2 IP Protocol

The value of the field specifies the matching value for the IP Protocol field [33]. 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
[22/23/60].9.2	2	prot1, prot2

Valid range: 0 to 257.

C.2.1.6.3 IPv4 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 $\text{src} = (\text{ip-src AND smask})$, where "smask" is the parameter from clause C.2.1.6.4. If this parameter is omitted, then comparison of the IP packet source address for this entry is irrelevant.

Type	Length	Value
[22/23/60].9.3	4	src1,src2,src3,src4

C.2.1.6.4 IPv4 Source Mask

The value of the field specifies the mask value for the IP source address, as described in clause C.2.1.6.3. If this parameter is omitted, then the default IP source mask is 255.255.255.255.

Type	Length	Value
[22/23/60].9.4	4	smask1,smask2,smask3,smask4

C.2.1.6.5 IPv4 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 $\text{dst} = (\text{ip-dst} \text{ AND } \text{dmask})$, where "dmask" is the parameter from clause C.2.1.6.6. If this parameter is omitted, then comparison of the IP packet destination address for this entry is irrelevant.

Type	Length	Value
[22/23/60].9.5	4	dst1,dst2,dst3,dst4

C.2.1.6.6 IPv4 Destination Mask

The value of the field specifies the mask value for the IP destination address, as described in clause C.2.1.6.5. If this parameter is omitted, then the default IP destination mask is 255.255.255.255.

Type	Length	Value
[22/23/60].9.6	4	dmask1,dmask2,dmask3,dmask4

C.2.1.7 TCP/UDP Packet Classification Encodings

This field defines the parameters associated with TCP/UDP packet classification.

While the TCP/UDP Packet Classification Encodings are located within the same subtype as the IPv4 Packet Classification Encodings, they apply regardless of IP version. The presence of an additional criterion from clause C.2.1.6 would cause the classifier to match only IPv4 packets. The presence of an additional criterion from clause C.2.1.10 would cause the classifier to match only IPv6 packets. For upstream classifiers, an Ethertype encoding indicating Ethertype 0x0800 or 0x86DD could also be used to cause the classifier to match only IPv4 or only IPv6 packets.

C.2.1.7.1 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 $\text{sportlow} \leq \text{src-port} \leq \text{sporthigh}$. 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
[22/23/60].9.7	2	sportlow1,sportlow2

C.2.1.7.2 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 $\text{sportlow} \leq \text{src-port} \leq \text{sporthigh}$. If this parameter is omitted, then the default value of sporthigh is 65 535. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23/60].9.8	2	sporthigh1,sporthigh2

C.2.1.7.3 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 $\text{dportlow} \leq \text{dst-port} \leq \text{dporhigh}$. 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
[22/23/60].9.9	2	dportlow1,dportlow2

C.2.1.7.4 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 $dportlow \leq dst-port \leq dporthigh$. If this parameter is omitted, then the default value of dporthigh is 65 535. This parameter is irrelevant for non-TCP/UDP IP traffic.

Type	Length	Value
[22/23/60].9.10	2	dporthigh1,dporthigh2

C.2.1.8 Ethernet LLC Packet Classification Encodings

This field defines the parameters associated with Ethernet LLC packet classification.

Type	Length	Value
[22/23/60].10	n	

C.2.1.8.1 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 $dst = (etherdst \text{ AND } msk)$. If this parameter is omitted, then comparison of the Ethernet MAC destination address for this entry is irrelevant.

Type	Length	Value
[22/23/60].10.1	12	dst1, dst2, dst3, dst4, dst5, dst6, msk1, msk2, msk3, msk4, msk5, msk6

C.2.1.8.2 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
[22/23/60].10.2	6	src1, src2, src3, src4, src5, src6

C.2.1.8.3 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 criterion. 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 [28] 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:

- Type 4: RNG-REQ.
- Type 6: REG-REQ.
- Type 7: REG-RSP.

- Type 14: REG-ACK.
- Type 30: INIT-RNG-REQ.
- Type 34: B-INIT-RNG-REQ.
- Type 44: REG-REQ-MP.
- Type 45: REG-RSP-MP.

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 after 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
[22/23/60].10.3	3	type, eprot1, eprot2

C.2.1.9 IEEE 802.1P/Q Packet Classification Encodings

This field defines the parameters associated with IEEE 802.1P/Q packet classification.

Type	Length	Value
[22/23/60].11	n	

C.2.1.9.1 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 $\text{pri-low} \leq \text{priority} \leq \text{pri-high}$. 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
[22/23/60].11.1	2	pri-low, pri-high

Valid Range is 0 to 7 for pri-low and pri-high.

C.2.1.9.2 IEEE 802.1Q VLAN_ID

The value of the field specifies 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
[22/23/60].11.2	2	vlan_id1, vlan_id2

C.2.1.10 IPv6 Packet Classification Encodings

This field defines the parameters associated with IPv6 packet classification. TCP/UDP Packet Classification Encodings (see clause C.2.1.7) are defined for IPv4 or IPv6 and may be present in a Service Flow Classifier of either type. If those classifiers are present in combination with IPv6 classifier encodings, then they apply to the IPv6 classifiers. If other IPv4 classifier encodings (22/23.9.1 thru 22/23.9.6) are present in the Service Flow Classifier along with IPv6 classifier encodings, then the Service Flow Classifier is invalid. If an invalid Service Flow Classifier of this type is sent to the CMTS in a Registration Request, the CMTS MUST reject the Registration Request. If an invalid Service Flow Classifier of this type is sent to the CMTS in a DSA or DSC Request message, the CMTS MUST reject the DSA or DSC Request message. If an invalid Service Flow Classifier of this type is sent to the CM in a DSA or DSC Request message, the CM MUST reject the DSA or DSC Request message.

Type	Length	Value
[22/23/60].12	n	

C.2.1.10.1 IPv6 Traffic Class Range and Mask

The values of the field specify the matching parameters for the IPv6 Traffic Class byte range and mask. An IP packet with IPv6 Traffic Class value "ip-tc" matches this parameter if $tc\text{-}low \leq (ip\text{-}tc \text{ AND } tc\text{-}mask) \leq tc\text{-}high$. If this field is omitted, then comparison of the IPv6 packet Traffic Class byte for this entry is irrelevant.

Type	Length	Value
[22/23/60].12.1	3	tc-low, tc-high, tc-mask

NOTE: The value 0x3F for tc-mask will exclude the Explicit Congestion Notification [i.5] bits from the comparison and hence will result in classification based on DSCP [43].

C.2.1.10.2 IPv6 Flow Label

The value of the field specifies the parameters of IPv6 flow label field in the IPv6 header. The 20 least significant bits represent the 20-bit IPv6 Flow Label while the 12 most significant bits are ignored. If this parameter is omitted, then comparison of IPv6 flow label for this entry is irrelevant.

Type	Length	Value
[22/23/60].12.2	4	FlowLabel

C.2.1.10.3 IPv6 Next Header Type

The value of the field specifies the desired upper-layer protocol type specified in the IPv6 header or extension headers associated with the packet. If this parameter is omitted, then comparison of any IPv6 next header type value for this entry is irrelevant.

The CM and CMTS MUST recognize the following Next Header types when searching for the upper-layer header.

Header Type	Description
0	Hop-by-Hop
60	Destination
43	Routing
44	Fragment
51	Authentication
50	Encapsulation
59	No

The CM and the CMTS look for the first Next Header field with a value that is not included in the above list in order to identify the Upper Layer protocol of the packet. The CMTS MUST apply the classifier rule to the packet according to the Upper Layer protocol that it identifies. The CM MUST apply the classifier rule to the packet according to the Upper Layer protocol that it identifies. If the CMTS initiates a transaction that configures a Classifier Rule with a Next Header value equal to one in the above list, the CM MUST reject that transaction. If the CM initiates a transaction that configures a Classifier Rule with a Next Header value equal to one in the above list, the CMTS MUST reject that transaction.

If a packet contains an ESP header, then it is assumed that the upper-layer header is encrypted and cannot be read. If the CM or CMTS encounters a packet with an ESP header then it MUST NOT match the packet to the Classifier Rule unless the classifier parameter value equals 256, as explained below.

If a packet is fragmented then a classifier might not be able to identify the upper-layer protocol of the second and following fragments.

There are two special IPv6 next header type field values: "256" that matches all IPv6 traffic, regardless of the Next Header values and "257" that matches both TCP and UDP traffic. An entry that includes an IPv6 next header type value greater than 257 MUST be invalidated for comparisons (i.e. no traffic can match this entry).

Type	Length	Value
[22/23/60].12.3	2	nhdr

C.2.1.10.4 IPv6 Source Address

The value of the field specifies the matching value for the IPv6 source address. An IPv6 packet with IPv6 source address "ip6-src" matches this parameter if $\text{src} = (\text{ip6-src AND smask})$. "smask" is computed by setting the most significant 'n' bits of smask to 1, where 'n' is IPv6 Source Prefix Length in bits. If the IPv6 Source Address parameter is omitted, then comparison of the IPv6 packet source address for this entry is irrelevant.

Type	Length	Value
[22/23/60].12.4	16	src

C.2.1.10.5 IPv6 Source Prefix Length (bits)

The value of the field specifies the fixed, most significant bits of an IPv6 address that are used to determine address range and subnet ID. If this parameter is omitted, then assume a default value of 128.

Type	Length	Value
[22/23/60].12.5	1	0 to 128

C.2.1.10.6 IPv6 Destination Address

The value of the field specifies the matching value for the IPv6 destination address. An IPv6 packet with IPv6 destination address "ip6-dst" matches this parameter if $\text{dst} = (\text{ip6-dst AND dmask})$. "dmask" is computed by setting the most significant 'n' bits of dmask to 1, where 'n' is IPv6 Destination Prefix Length in bits. If the IPv6 Destination Address parameter is omitted, then comparison of the IPv6 packet destination address for this entry is irrelevant.

Type	Length	Value
[22/23/60].12.6	16	dst

C.2.1.10.7 IPv6 Destination Prefix Length (bits)

The value of the field specifies the fixed, most significant bits of an IPv6 address that are used to determine address range and subnet ID. If this parameter is omitted, then assume a default value of 128.

Type	Length	Value
[22/23/60].12.7	1	0 to 128

C.2.1.11 Vendor Specific Classifier Parameters

This allows vendors to encode vendor-specific classifier parameters using the DOCSIS Extension Field. 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 clause C.1.1.17).

Type	Length	Value
[22/23/60].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 composed from a number of encapsulated type/length/value fields.

NOTE: 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 composed from a number of encapsulated type/length/value fields.

NOTE: 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 to 65 535

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 to 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 messages related to establishing an admitted or active upstream Service Flow. This field MUST also be present in REG-RSP, REG-RSP-MP, 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 signalling 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-4: Values Used in REG-REQ, REG-REQ-MP, REG-RSP and REG-RSP-MP 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-5: Values Used In REG-REQ, REG-REQ-MP, REG-RSP, REG-RSP-MP 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 clause 7.5.7.4).

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 (see clause C.4).

C.2.2.3.6 Service Flow Required Attribute Mask

This parameter is optional in upstream and downstream service flows. If specified, it limits the set of channels and bonding groups to which the CMTS assigns the service flow requiring certain Cable Operator-determined binary attributes. When this TLV is not present in the service flow request the CMTS defaults this value to zero.

Type	Length	Value
[24/25].31	4	32-bit mask representing the set of binary channel attributes required for service flow. This TLV uses the BITS Encoding convention where bit number 0 is the most significant bit of the mask

See clause 8.1.1 for how the Service Flow Required Attribute mask, Service Flow Forbidden Attribute Mask and Service Flow Attribute Aggregation Rule Mask control how service flows may be assigned to particular channels or bonding groups.

C.2.2.3.7 Service Flow Forbidden Attribute Mask

This parameter is optional in upstream and downstream service flows. If specified, it limits the set of channels and bonding groups to which the CMTS assigns the service flow by forbidding certain attributes. When this TLV is not present in the service flow request the CMTS defaults this value to zero.

Type	Length	Value
[24/25].32	4	32-bit mask representing the set of binary channel attributes forbidden for the service flow. This TLV uses the BITS Encoding convention where bit number 0 is the most significant bit of the mask.

See clause 8.1.1 for how the Service Flow Required Attribute mask, Service Flow Forbidden Attribute Mask and Service Flow Attribute Aggregation Rule Mask control how service flows may be assigned to particular channels or bonding groups.

C.2.2.3.8 Service Flow Attribute Aggregation Rule Mask

This parameter is optional in upstream and downstream service flows. It controls, on a per-attribute basis, whether the attribute is required or forbidden on any or all channels of a bonding group that aggregates multiple channels. It can be considered to control how an "aggregate" attribute mask for the bonding group is built by either AND'ing or OR'ing the attributes of individual channels of the bonding group. When this TLV is not present in the service flow request the CMTS defaults this value to zero.

Type	Length	Value
[24/25].33	4	32-bit mask controlling how attributes in each bit position are aggregated for bonding groups consisting of multiple channels. A '1' in this mask for an attribute means that a bonding group attribute is considered to be the logical 'AND' of the attribute bit for each channel. A '0' in this mask for an attribute means that the bonding group is considered to have the logical 'OR' of the attribute for each channel. This TLV uses the BITS Encoding convention where bit number 0 is the most significant bit of the mask.

See clause 8.1.1 for how the Service Flow Required Attribute mask, Service Flow Forbidden Attribute Mask and Service Flow Attribute Aggregation Rule Mask control how service flows may be assigned to particular channels or bonding groups.

C.2.2.3.9 Application Identifier

This parameter allows for the configuration of a Cable Operator defined Application Identifier for service flows, e.g. an Application Manager ID and Application Type as defined in [23]. This Application Identifier can be used to influence admission control or other policies in the CMTS that are outside of the scope of the present document.

Type	Length	Value
[24/25].34	4	Application ID

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, REG-RSP-MP, 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, REG-REQ-MP, 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, REG-RSP-MP, DSA-RSP or DSC-RSP.

On failure, the REG-RSP, REG-RSP-MP, DSA-RSP or DSC-RSP MUST include one Service Flow Error Encoding for at least one failed Service Flow requested in the REG-REQ, REG-REQ-MP, 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, REG-RSP-MP, DSA-RSP or DSC-RSP message. A Service Flow Error Encoding for the failed Service Flow MUST include the Confirmation Code and Errored Parameter. A Service Flow Error Encoding for the failed Service Flow 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, REG-RSP-MP, 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, REG-REQ-MP, 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 clause 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.
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 and 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. The CMTS **SHOULD** provide differentiated service based on the value of Traffic Priority. The specific algorithm for enforcing this parameter is not mandated here. The default priority is 0.

For upstream service flows, the CMTS **SHOULD** use this parameter when determining precedence in request service and grant generation. For upstream service flows, the CM **MUST** include contention Request opportunities for Priority Request Service IDs (refer to clause A.2.3) in its request backoff algorithm based on this priority and its Request/Transmission Policy (refer to clause C.2.2.6.3).

For downstream service flows configured with a non-default value, the CMTS inserts this priority as a three bit tag into the Downstream Service Extended Header as defined in clause 6.2.5.6. The CM preferentially orders the PDU packets onto the egress queues based on this 3-bit Traffic Priority in the DS EHDR as described in clause 7.6.

Type	Length	Value
[24/25].7	1	0 to 7 - Higher numbers indicate higher priority

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, including every PDU in the case of a Concatenated MAC Frame. This parameter is applied after Payload Header Suppression; it does not include the bytes suppressed for PHS.

The number of bytes forwarded (in bytes) is limited during any time interval T by Max(T), as described in the expression:

$$\text{Max}(T) = T * (R / 8) + B \quad (1)$$

where the parameter B (in bytes) is the Maximum Traffic Burst Configuration Setting (refer to clause C.2.2.5.3).

NOTE 1: This parameter does not limit the instantaneous rate of the Service Flow.

The specific algorithm for enforcing this parameter is not mandated here. Any implementation which satisfies the above equation is conformant. In particular, the granularity of enforcement and the minimum implemented value of this parameter are vendor specific. The CMTS SHOULD support a granularity of at most 100 kbps. The CM SHOULD support a granularity of at most 100 kbps.

NOTE 2: 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 expression (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 expression (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 expression (1) in any interval T. The CMTS SHOULD "rate shape" the downstream traffic by enqueueing packets arriving in excess of expression (1) and delay them until the expression can be met.

When a CMTS implements both a Maximum Sustained Traffic Rate and a Peak Downstream Traffic Rate for a service flow, it limits the bytes forwarded in any interval T to the lesser of Max(T) defined in equation (1) and Peak(T) defined in equation (2) of clause C.2.2.7.2.

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, including every PDU in the case of a Concatenated MAC Frame.

The minimum value of B is 1 522 bytes. If this parameter is omitted, the default value for B is 3 044 bytes. This parameter has no effect unless a non-zero value has been provided for the Maximum Sustained Traffic Rate parameter.

Bonded downstream packets may be internally distributed across multiple channels within the CMTS after they have been scheduled according to the rate limiting algorithm in expression (1). As a result, the traffic burst observed at the CMTS output would not just be a function of the rate limiting algorithm, but would also be a function of the skew between the channels that data is sent on. Thus the observed traffic burst could exceed the Maximum Traffic Burst value.

The resequencing and reassembly operations may also impact the observed maximum traffic burst of a downstream or upstream bonded service flow. When a stream of packets are resequenced (or segments are reassembled) they can not be forwarded until all have arrived (or a timeout occurred). As a result, a period of idle time would be followed by a traffic burst even if the CMTS/CM performed perfect output shaping of the traffic as per (1).

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 value of this parameter affects the trade-off 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 burst 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 value of this parameter is calculated from the byte following the MAC header HCS to the end of the CRC, including every PDU in a Concatenated MAC Frame. If this parameter is omitted, then it defaults to a value of 0 bit/sec (i.e. no bandwidth is reserved for the flow by default).

How Minimum Reserved Traffic Rate and Assumed Minimum Reserved Rate Packet Size apply to a CMTS's admission control policies is vendor specific and is beyond the scope of the present document. The aggregate Minimum Reserved Traffic Rate of all Service Flows could exceed the amount of available bandwidth.

Unless explicitly configured otherwise, a CMTS SHOULD schedule forwarding of all service flows' traffic such that each receives at least its Minimum Reserved Traffic Rate when transmitting packets with the Assumed Minimum Reserved Rate Packet Size. If the service flow sends packets of a size smaller than the Assumed Minimum Reserved Rate Packet Size, such packets will be treated as being of the Assumed Minimum Reserved Rate Packet Size for calculating the rate forwarded from the service flow for purposes of meeting the 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.

NOTE: The granularity of the Minimum Reserved Traffic Rate used internally by the CMTS is vendor specific. Because of this, the CMTS MAY schedule forwarding of a service flow's traffic at a rate greater than the configured value for Minimum Reserved Traffic Rate.

This field is only applicable at the CMTS.

Type	Length	Value
[24/25].10	4	

C.2.2.5.5 Assumed Minimum Reserved Rate Packet Size

This parameter is used by the CMTS to make worst-case DOCSIS overhead assumptions. The Minimum Reserved Traffic Rate of a service flow excludes the DOCSIS MAC header and any other DOCSIS overhead (e.g. for completing an upstream mini-slot). Traffic with smaller packets sizes will require a higher proportion of overall channel capacity for DOCSIS overhead than traffic with larger packet sizes. The CMTS assumes that the worst-case DOCSIS overhead for a service flow will be when all traffic is as small as the size specified in this parameter.

This parameter is defined in bytes and is specified as the bytes following the MAC header HCS to the end of the CRC.

If this parameter is omitted, then the default value is CMTS implementation dependent.

Type	Length	Value
[24/25].11	2	

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. This parameter 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 QoS 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 clause 7.5.5.2), 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. This parameter **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 s 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 using the DOCSIS Extension Field. 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 clause C.1.1.17.)

Type	Length	Value
[24/25].43	N	

C.2.2.5.9 IP Type Of Service (DSCP) Overwrite

The CMTS **MUST** overwrite IP packets with IPv4 TOS byte or IPv6 Traffic Class 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/Traffic Class byte is not overwritten.

This parameter is only applicable at the CMTS. If defined, this parameter **MUST** be enforced by the CMTS.

The IPv4 TOS octet as originally defined in RFC 791 [i.9] has been superseded by the 6-bit Differentiated Services Field (DSField, [i.6]) and the 2-bit Explicit Congestion Notification Field (ECN field, [i.5]). The IPv6 Traffic Class octet [41] is consistent with that new definition. Network operators should avoid specifying values of tos-and-mask and tos-or-mask that would result in the modification of the ECN bits.

In particular, operators should not use values of tos-and-mask that have either of the least-significant two bits set to 0. Similarly, operators should not use values of tos-or-mask that have either of the least-significant two bits set to 1.

Type	Length	Value
24/25.23	2	tos-and-mask, tos-or-mask

C.2.2.5.10 Peak Traffic Rate

This parameter is the rate parameter P of a token-bucket-based peak rate limiter for packets of a service flow. Configuring this peak rate parameter permits an operator to define a Maximum Traffic Burst value for the Maximum Sustained Traffic Rate much larger than a maximum packet size, but still limit the burst of packets consecutively transmitted for a service flow (refer to clause C.2.2.5.3).

The parameter P is expressed in bits per second and includes all MAC frame data PDU bytes scheduled on the service flow from the byte following the MAC header HCS to the end of the CRC. This parameter is applied after Payload Header Suppression; it does not include the bytes suppressed for PHS.

The number of bytes forwarded is limited during any time interval T by Peak(T), as described in expression (2), below:

$$\text{Peak}(T) \leq T * (P / 8) + 1 \quad 522 \quad (2)$$

C.2.2.5.10.1 Upstream Peak Traffic Rate

For an upstream Service Flow, the CM **SHOULD NOT** request bandwidth exceeding the Peak(T) requirement in expression (2) during any interval T because this could force the CMTS to discard packets and/or fill MAPs with deferred grants.

The CM SHOULD defer upstream packets that violate expression (2) and "rate shape" them to meet the expression, up to a limit as implemented by vendor buffering restrictions.

The CMTS SHOULD enforce expression (2) 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 SHOULD 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
24.27	4	Upstream Peak Traffic Rate (P), in bits per second. If omitted or zero(0), upstream peak traffic rate is not limited

C.2.2.5.10.2 Downstream Peak Traffic Rate

When this parameter P is defined for a service flow, the CMTS SHOULD enforce the number of PDU bytes scheduled on a downstream service flow for any time interval T to be limited by the expression Peak(T) as described in expression (2).

When a CMTS implements both a Maximum Sustained Traffic Rate and a Peak Downstream Traffic Rate for a service flow, it limits the bytes forwarded in any interval T to the lesser of Max(T) defined in equation (1) of clause C.2.2.5.2 and Peak(T) defined in equation (2). The peak rate parameter P is intended to be configured to be greater than or equal to the Maximum Sustained Rate R of equation (1). Operation when the peak rate P is configured to be less than the Maximum Sustained Rate R is CMTS vendor-specific.

When the CMTS enforces the Downstream Peak Traffic Rate, it SHOULD "rate shape" the downstream traffic by delaying the forwarding of packets until the Downstream Peak Rate expression (2) can be met. The specific algorithm for enforcing this parameter, with or without concurrently enforcing the Maximum Sustained Traffic Rate parameter, is not mandated here. Any implementation which satisfies the normative requirements is conformant. In particular, the granularity of enforcement and the minimum implemented value of this parameter are vendor specific. The CMTS SHOULD support a granularity of at most 100 kbps.

This parameter is not intended for enforcement on the CM.

If the parameter is omitted or set to zero, the CMTS MUST NOT enforce a Downstream Peak Traffic Rate for the service flow.

Type	Length	Value
25.27	4	Downstream Peak Traffic Rate (P), in bits per second. If omitted or zero(0), downstream peak traffic rate is not limited.

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 when not operating in MTC Mode. 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 1 522.

This field is only applicable at the CM. If defined, this parameter MUST be enforced at the CM.

NOTE 0: This value does not include any physical layer overhead.

Type	Length	Value
24.14	2	

NOTE 1: This applies only to concatenated bursts and only when the CM is not operating in MTC Mode. 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 2: 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 - see note)
		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
NOTE:	The specific implementation dependent scheduling service type could be defined in the 24.43 Vendor Specific QoS Parameters. (Refer to clause C.2.2.5.8.)	

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 clause 7.2.3 for requirements related to settings of the bits of this parameter for each Service Flow Scheduling Type. For Continuous Concatenation and Fragmentation, it specifies whether or not segment headers are used and what opportunities can be used for making bandwidth requests.

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 clause 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 (see note 1).
		Bit #4 The Service Flow MUST NOT piggyback requests with data.
		Bit #5 The Service Flow MUST NOT concatenate data (see note 2).
		Bit #6 The Service Flow MUST NOT fragment data.
		Bit #7 The Service Flow MUST NOT suppress payload headers.
		Bit #8 The Service Flow MUST drop packets that do not fit in the Unsolicited Grant Size (see notes 3 and 4).
		Bit #9 The Service Flow MUST NOT use segment headers. When set to zero, the Service Flow MUST use segment headers (see note 5).
		Bit #10 The Service Flow MUST NOT use contention regions for transmitting multiple outstanding bandwidth requests.
		All other bits are reserved.
NOTE 1: This bit is irrelevant for a CM in Multiple Transmit Channel Mode because it does not use Request/Data for sending data.		
NOTE 2: This bit applies for pre-3.0 DOCSIS operation.		
NOTE 3: 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.		
NOTE 4: 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.		
NOTE 5: Only UGS or UGS-AD Service Flows can be configured with Segment Header OFF for CMs operating in Multiple Transmit Channel Mode. 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 $t_i = t_0 + i \cdot \text{interval}$. The actual poll times, t'_i MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, 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 clause 7.1).

This field is only applicable at the CMTS. If defined, this parameter MUST be enforced by the CMTS.

Type	Length	Value
24.17	4	Number of microseconds

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 $t_i = t_0 + i \cdot \text{interval}$. The actual poll, t'_i MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, 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 clause 7.1).

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	Number of microseconds

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 beginning with the Frame Control byte for Segment Header OFF operation or the first byte of the Segment Header for Segment Header ON operation and ending at the 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	Number of bytes

NOTE: For UGS, this parameter should be used by the CMTS to compute the size of the unsolicited grant in mini-slots.

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 $t_i = t_0 + i \cdot \text{interval}$. The actual grant times, t'_i MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, 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 are 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 clause 7.1).

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	Number of microseconds

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 $t_i = t_0 + i \cdot \text{interval}$. The actual transmission opportunities, t_i MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, 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 clause 7.1).

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	Number of microseconds

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 $t_i = t_0 + i \cdot \text{interval}$. The actual grant times, t'_i MUST be in the range $t_i \leq t'_i \leq t_i + \text{jitter}$, 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 are 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
24.22	1	# of grants (valid range: 0 to 127)

C.2.2.6.10 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 $t_i = t_0 + i \cdot \text{interval}$, where interval is the Nominal Grant Interval (refer to clause 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
24.24	4	CMTS Timestamp (valid range: 0 to 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:

- $\text{interval} - (((\text{current time} - t_0) / 10.24) \text{ modulus interval});$

where interval is in units of microseconds and current time and t_0 are in 10,24 MHz units.

C.2.2.6.11 Multiplier to Contention Request Backoff Window

In 3.0 operation, this is a multiplier to be applied by a CM performing contention request backoff for data requests. Clause 7.2.1.5 contains the details on how this multiplier is applied. This setting is not included in a CM configuration file. The CMTS MAY include this setting whenever it provides a CM the parameters associated with a service flow.

Type	Length	Value
24.25	1	Number of eighths (valid range: 4-12)

If this parameter is not encoded, the parameter value is assumed to be 8 and thus, the multiplier is equal to 1. If the received value is outside the valid range, the CM MUST assume a value of 8 and thus, the multiplier is equal to 1.

C.2.2.6.12 Multiplier to Number of Bytes Requested

In 3.0 operation, this is a multiplier to be assumed in any bandwidth request (REQ burst or piggyback request). Clause 7.2.1.4 contains the details on how this multiplier is applied.

Type	Length	Value
24.26	1	Multiplying factor (valid range: 1, 2, 4, 8 or 16)

If this parameter is not encoded, the default value of 4 is used.

C.2.2.7 Downstream-Specific QoS Parameter Encodings

C.2.2.7.1 Maximum Downstream Latency

The value of this parameter specifies the desired maximum latency across the DOCSIS network, beginning with the reception of a packet by the CMTS on its NSI and including the transit of the CIN (if applicable), the forwarding of the packet on an RF Interface and (in the case of sequenced traffic) the release of the packet from the Resequencing operation in the CM.

This parameter is intended to influence the CMTS scheduling, M-CMTS DEPI flow assignment and assignment of the service flow to downstream bonding groups. The CMTS SHOULD attempt to meet the desired maximum downstream latency.

When this parameter is defined, the CMTS MUST NOT transmit the packets of the Service Flow using a Resequencing DSID that has a Max_Resequencing_Wait in excess of the value of this parameter.

Type	Length	Value
25.14	4	Number of microseconds

The value of 0 is equivalent to the TLV not present, i.e. no limitations on latency specified.

C.2.2.7.2 Downstream Resequencing

This parameter controls resequencing for downstream service flows. In particular, this parameter controls whether or not the service flow is to be associated with a Resequencing DSID. When a service flow is associated with a Resequencing DSID, a sequence number is inserted in the 5-byte DS EHDR on every packet. See clauses 6.2.5.6 and 8.2.3.

Type	Length	Value
25.17	1	0 = The CMTS MUST associate this service flow with a resequencing DSID if the service flow is assigned to a downstream bonding group. 1 = The CMTS MUST NOT associate this service flow with a resequencing DSID.

If this TLV is not present, a default value of 0 MUST be used by the CMTS.

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 clause C.2.1.4.1).

Type	Length	Value
26.1	1	1 to 255

C.2.2.8.2 Classifier Identifier

The value of the field specifies a Classifier Identifier that identifies the corresponding Classifier (refer to clause C.2.1.4.2).

Type	Length	Value
26.2	2	1 to 65 535

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 clause C.2.1.4.3).

Type	Length	Value
26.3	2	1 to 65 535

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 to 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

For PHSI-indexed PHS, 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 all PHS Rules" command is used to delete all PHS Rules for a specified Service Flow. See clause 6.4.15 for details on DSC-REQ required PHS parameters when using this option.

NOTE: An attempt to add a PHS Rule which already exists is an error condition. An attempt to delete a PHS Rule which does not exist is also an error condition.

C.2.2.8.6 Dynamic Bonding Change Action

When received in a Dynamic Bonding Change Request, this indicates the action that MUST be taken with this payload header suppression byte string.

Type	Length	Value
26.13	1	0 - Add PHS Rule
		1 - Delete PHS Rule

For DSID-indexed PHS, the only valid actions are "Add PHS Rule" and "Delete PHS Rule".

NOTE: An attempt to add a PHS Rule which already exists is an error condition. An attempt to delete a PHS Rule which does not exist is also 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, REG-RSP-MP, DSA-RSP, DSC-RSP and DBC-RSP messages to indicate the reason for the recipient's negative response to a Payload Header Suppression Rule establishment request in a REG-REQ, REG-REQ-MP, DSA-REQ, DSC-REQ or DBC-REQ message.

On failure, the REG-RSP, REG-RSP-MP, DSA-RSP, DSC-RSP or DBC-RSP MUST include one Payload Header Suppression Error Encoding for at least one failed Payload Header Suppression Rule requested in the REG-REQ, REG-REQ-MP, DSA-REQ, DSC-REQ or DBC-REQ message. A Payload Header Suppression Error Encoding for the failed Payload Header Suppression Rule MUST include the Confirmation Code and Errored Parameter. A Payload Header Suppression Error Encoding for the failed Payload Header Suppression Rule 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, REG-RSP-MP, DSA-RSP, DSC-RSP or DBC-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 valid REG-REQ, REG-REQ-MP, DSA-REQ, DSC-REQ or DBC-REQ message does not contain a Payload Header Suppression Error Encoding.

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 clause 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. 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 contents 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 implementation 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 254 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 = do not suppress first byte of the suppression field 1 = suppress first byte of the suppression field
		bit 1:0 = do not suppress second byte of the suppression field 1 = suppress second byte of the suppression field
		bit x:0 = do not 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. If the bit value is a "1" (and verification passes or is disabled), 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. If the bit value is a "0", 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 byte, 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 = do not 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 clause 7.7.3.)

C.2.2.10.6 Vendor Specific PHS Parameters

This allows vendors to encode vendor-specific PHS parameters using the DOCSIS Extension Field. 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 clause C.1.1.7.)

Type	Length	Value
26.43	n	

C.3 Encodings for Other Interfaces

C.3.1 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 [15].

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 Bonding Change-Response, 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 table C-6 are used both as message Confirmation Codes and as Error Codes in Error Set Encodings which may be carried in these messages.

Confirmation codes 200 to 220 are reserved for Major Errors. These confirmation codes MUST be used only as message Confirmation Codes. 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 or to generate a full RSP message.

Table C-6: Confirmation Codes

Confirmation	Conf. code	Description	Applicable Message(s)								
			REG-RSP	REG-ACK	DSA-RSP	DSA-ACK	DSC-RSP	DSC-ACK	DSD-RSP	DCC-RSP	DBC-RSP
okay / success	0	the message was received and successful.	X	X	X	X	X	X	X	X	X
reject-other	1	none of the other reason codes apply.	X	X	X	X	X	X	X	X	X
reject-unrecognized-configuration-setting	2	a configuration setting or TLV value is outside of the specified range.	X	X	X	X	X	X	X	X	X
reject-temporary / reject-resource	3	the current loading of the CMTS or CM prevents granting the request, but the request might succeed at another time.	X	X	X	X	X	X	X	X	X
reject-permanent / reject-admin	4	for policy, configuration or capabilities reasons, the request would never be granted unless the CMTS or CM were manually reconfigured or replaced.	X	X	X	X	X	X	X	X	X
reject-not-owner	5	the requester is not associated with this service flow.	X	X	X	X	X	X	X	X	X
reject-service-flow-not-found	6	the Service Flow indicated in the request does not exist.	X	X	X	X	X	X	X	X	X
reject-service-flow-exists	7	the Service Flow to be added already exists.			X						
reject-required-parameter-not-present	8	a required parameter has been omitted.	X	X	X	X	X	X	X	X	X
reject-header-suppression	9	the requested header suppression cannot be supported.	X	X	X	X	X	X			X
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).				X		X			
reject-authentication-failure	11	the requested transaction was rejected because the message contained an invalid HMAC-digest, CMTS-MIC, provisioned IP address or timestamp.	X	X	X	X	X	X	X	X	X
reject-add-aborted	12	the addition of a dynamic service flow was aborted by the initiator of the Dynamic Service Addition.				X					
reject-multiple-errors	13	multiple errors have been detected.	X	X	X	X	X	X	X	X	X
reject-classifier-not-found	14	the request contains an unrecognized classifier ID.			X	X	X	X			
reject-classifier-exists	15	the ID of a classifier to be added already exists.			X	X	X	X			
reject-PHS-rule-not-found	16	the request references a PHS rule that does not exist.					X				X
reject-PHS-rule-exists	17	the request attempts to add a PHS rule that already exists.			X		X				X
reject-duplicate-reference-ID-or-index-in-message	18	the request used a service flow reference, classifier reference, SFID, DSID, SAID or classifier ID twice in an illegal way.	X	X	X	X	X	X	X	X	X
reject-multiple-upstream-service-flows	19	DSA/DSC/DSD contains parameters for more than one upstream flow.			X	X	X	X	X		
reject-multiple-downstream-service-flows	20	DSA/DSC/DSD contains parameters for more than one downstream flow.			X	X	X	X	X		
reject-classifier-for-another-service-flow	21	DSA/DSC-REQ includes classifier parameters for a SF other than the SF(s) being added/changed by the DSA/DSC.			X		X				
reject-PHS-for-another-service-flow	22	DSA/DSC-REQ includes a PHS rule for a SF other than the SF(s) being added/changed by the DSA/DSC.			X		X				
reject-parameter-invalid-for-context	23	the parameter supplied cannot be used in the encoding in which it was included or the value of a parameter is invalid for the encoding in which it was included.	X	X	X	X	X	X	X	X	X
reject-authorization-failure	24	the requested transaction was rejected by the authorization module.	X		X		X				

Confirmation	Conf. code	Description	Applicable Message(s)								
			REG-RSP	REG-ACK	DSA-RSP	DSA-ACK	DSC-RSP	DSC-ACK	DSD-RSP	DCC-RSP	DBC-RSP
reject-temporary-DCC	25	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.			X		X				
reject-downstream-inconsistency	26	the RCS and DS Resequencing Channel Lists are inconsistent.		X							X
reject-upstream-inconsistency	27	the TCS and Service Flow SID Cluster assignments are inconsistent.		X	X	X	X	X			X
reject-insufficient-SID-cluster-resources	28	the SID Cluster assignment would require more SID Clusters than the CM has available.		X	X	X	X	X			X
reject-missing-RCP	29	there was no RCP included with the modem's registration request, although it indicated support for Multiple Receive Channel Mode.	X								
partial-service	30	CM unable to use one or more channels as instructed in the DBC-REQ or REG-RSP.		X							X
reject-temporary-DBC	31	CMTS needs to perform a DBC in order to execute a DSA or DSC.			X		X				
reject-unknown-DSID	32	DBC-REQ trying to change attributes of an unknown DSID.									X
reject-unknown-SID-Cluster	33	Unknown SID Cluster ID.									X
reject-invalid-initialization-technique	34	Initialization technique not permitted or not within the values known to the CM.		X						X	X
reject-no-change	35	CM is already using all the parameters specified in the DBC-REQ.									X
reject-invalid-DBC-request	36	CM is rejecting DBC-REQ as invalid, per clause 11.5.2.									X
reject-mode-switch	37	DBC-REQ requires CM to switch from legacy mode to Multiple Transmit Channel Mode.									X
reject-insufficient-transmitters	38	implementation would require more upstream transmitters than the CM has available.		X							X
reject-insufficient-DSID-resources	40	implementation would require more DSIDs than the CM has available.		X							X
reject-invalid-DSID-encoding	41	The message has an invalid DSID encoding.		X							X
reject-unknown-client-mac-address	42	DSID Multicast Client MAC address is not known by the CM.		X							X
reject-unknown-SAID	43	The message attempts to delete an unknown SAID.									X
reject-insufficient-SA-resources	44	implementation would require more SAIDs than the CM has available.		X							X
reject-invalid-SA-encoding	45	The message has an invalid SA encoding.		X							X
reject-invalid-SA-crypto-suite	46	The message has an invalid SA crypto suite.		X							X
reject-tek-exists	47	CMTS attempts to set an SA at the CM for which the CM already has an active TEK state machine.		X							X
reject-invalid-SID-cluster-encoding	48	The message has an invalid SID cluster encoding.		X	X	X					X
reject-insufficient-SID-resources	49	The SID assignment would require more SIDs than the CM has available.		X	X	X					X
reject-unsupported-parameter-change	50	The DSC-REQ contains a parameter to be changed where the support for the change is optional and this device does not support it.					X				
reject-phs-rule-fully-defined	51	An attempt has been made in a DSC-REQ to Set a PHS element in a rule that is already fully defined.					X				
reject-NoMAPsOrUCDs	52	No MAPs or UCDs for the designated upstream channel.		X							X

Confirmation	Conf. code	Description	Applicable Message(s)								
			REG-RSP	REG-ACK	DSA-RSP	DSA-ACK	DSC-RSP	DSC-ACK	DSD-RSP	DCC-RSP	DBC-RSP
error-T3RetriesExceeded	53	16 consecutive T3 timeouts while trying to range on designate upstream channel.		X							X
error-T2Timeout	54	CM experienced T2 timeout on the designated upstream channel.		X							X
error-T4Timeout	55	CM experienced T4 timeout on the designated upstream channel.		X							X
error-RangeAbort	56	CM received RNG-RSP with Status ABORT on the designated upstream channel.		X							X
error-InitChanTimeout	57	Initializing Channel Timeout occurred before acquiring all channels.		X							X
error-DBC-REQ-incomplete	58	"DBC-REQ Timeout" timer expired before all fragments of the DBC-REQ message have been correctly received.									X
L2VPN-specific	100-109	These confirmation codes are reserved for L2VPN usage. See [8].									
reject-unknown-RCP-ID	160	RCP-ID in RCC not supported by CM.		X							X
reject-multiple-RCP-IDs	161	only one RCP-ID is allowed in RCC.		X							X
reject-missing-Receive-Module-Index	162	Receive Module Index missing in RCC.		X							X
reject-invalid-Receive-Module-Index	163	RCC contains a Receive Module Index which is not supported by CM.		X							X
reject-invalid-receive-channel-center-frequency	164	receive channel center frequency not within allowed range of center frequencies for Receive Module.		X							X
reject-invalid-RM-first-channel-center-frequency	165	Receive Module first channel center frequency not within allowed range of center frequencies.		X							X
reject-missing-RM-first-channel-center-frequency	166	Receive Module first channel center frequency not present in RCC.		X							X
reject-no-primary-downstream-channel-assigned	167	no primary downstream channel assignment in RCC.		X							X
reject-multiple-primary-downstream-channel-assigned	168	more than one primary downstream channel assignment present in RCC.		X							X
reject-receive-module-connectivity-error	169	Receive Module connectivity encoding in RCC requires configuration not supported by CM.		X							X
reject-invalid-receive-channel-index	170	receive channel index in RCC not supported by CM.		X							X
reject-center-frequency-not-multiple-of-62 500-Hz	171	center frequency in RCC not a multiple of 62 500 Hz.		X							X
depart	180	the CM is on the old channel and is about to perform the jump to the new channel.								X	
arrive	181	the CM has performed the jump and has arrived at the new channel.								X	
reject-already-there	182	the CMTS has asked the CM to move to a channel that it is already occupying as described in clause 11.4.2 or sent a DBC-REQ with redundant parameters as described in clause 11.5.2.								X	X
reject-20-disable	183	the CMTS has asked a CM with 2.0 mode disabled to move to a Type 3 channel that it cannot use and a UCD substitution was sent in the corresponding DCC-REQ.								X	
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.	X	X	X	X	X	X			X

Confirmation	Conf. code	Description	Applicable Message(s)								
			REG-RSP	REG-ACK	DSA-RSP	DSA-ACK	DSC-RSP	DSC-ACK	DSD-RSP	DCC-RSP	DBC-RSP
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.	X	X	X	X	X	X			X
reject-major-PHS-rule-error	202	indicates that the REQ message did not have both a Service Flow Reference/Identifier and a Classifier Reference/Identifier and that PHS rule major errors were the only major errors.	X	X	X	X	X	X			X
reject-multiple-major-errors	203	indicates that the REQ message contained multiple major errors of types 200, 201 or 202.	X	X	X	X	X	X			X
reject-message-syntax-error	204	indicates that the REQ message contained syntax error(s) (e.g. a TLV length error) resulting in parsing failure.	X	X	X	X	X	X	X	X	X
reject-primary-service-flow-error	205	indicates that a REG-REQ REG-REQ-MP, REG-RSP or REG-RSP-MP message did not define a required primary Service Flow or a required primary Service Flow was not specified active.	X	X							
reject-message-too-big	206	the length of the message needed to respond exceeds the maximum allowed message size.	X	X	X	X	X	X	X	X	X
reject-invalid-modem-capabilities	207	the REG-REQ or REG-REQ-MP contained either an invalid combination of modem capabilities or modem capabilities that are inconsistent with the services in the REG-REQ or REG-REQ-MP.	X								
reject-bad-rcc	208	the message contained an invalid Receive Channel Configuration.		X							X
reject-bad-tcc	209	the message contained an invalid Transmit Channel Configuration.		X							X
reject-dynamic-range-window-violation	210	channels added or deleted by the REG-RSP-MP or DBC-REQ would have resulted in a dynamic range window violation.		X							X

Annex D (normative): CM Configuration Interface Specification

D.1 CM Configuration

D.1.1 CM Binary Configuration File Format

The CM-specific configuration data are 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 [38].

It consists of a number of configuration settings (1 per parameter) each of the form: Type Length Value.

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

The CMs **MUST** support a 8 192-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 calculated 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.1.2 Configuration File Settings

The following configuration settings are included in the configuration file and **MUST** be supported by all CMs. The CM **MUST NOT** send a REG-REQ or REG-REQ-MP 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 1: A DOCSIS 1.0 CM needs to 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.
- Enable 2.0 Mode.
- Enable Test Modes.
- Static Multicast MAC Address.

The following configuration settings may be included in the configuration file; and if present, MAY be supported by a CM:

- DOCSIS Extension Field Configuration Settings.

NOTE 2: There is a limit on the size of Registration Request and Registration Response frames (see clause 8.2.5). The configuration file should not be so large as to require the CM or CMTS to exceed that limit. If the Extended CMTS MIC Encoding is included in the CM Configuration file, the CM MUST include in its REG-REQ or REG-REQ-MP message all instances of top-level TLVs in the CM configuration for which there is a '1' bit in the CMTS MIC Encoding Bitmask.

D.1.3 Configuration File Creation

The sequence of operations required to create the configuration file is as shown in figure D-2 through figure D-6.

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

type, length, value for parameter 1
type, length, value for parameter 2
type, length, value for parameter n

Figure D-2: Create TLV Entries for Parameters Required by the CM

- 2) Insert the Extended CMTS MIC Parameters configuration setting as defined in clause D.2.1 and add to the file following the last parameter using code and length values defined for this field. A configuration file for a pre-DOCSIS 3.0 modem MAY include the Extended CMTS MIC.

NOTE 1: The Extended CMTS MIC Encoding may include an Explicit Extended CMTS MIC Digest subtype that is calculated over the top-level parameters in the Extended CMTS MIC Bitmap ordered first by top-level TLV type code and secondly by their position within the CM configuration file (and hence their position in REG-REQ/REG-REQ-MP).

NOTE 2: The Explicit Extended CMTS MIC Digest value, if present, does not include either the CM MIC or CMTS MIC digest value. If the Explicit Extended CMTS MIC Digest value is present and the Extended CMTS MIC Bitmap indicates that TLV 43 is to be covered by the Extended CMTS MIC, then the Explicit Extended CMTS MIC Digest TLV is initially populated with an all-zeros value (and length appropriate for the HMAC Type selected) for purposes of the Extended CMTS MIC Digest calculation.

type, length, value for parameter 1
type, length, value for parameter 2
type, length, value for parameter n
type, length, value for Ext CMTS MIC Params

Figure D-3: Add Extended CMTS MIC Parameters

- 3) Calculate the CM message integrity check (MIC) configuration setting as defined in clause D.1.3.1 and add to the file following the Extended CMTS MIC Params using code and length values defined for this field.

NOTE 3: The CM MIC code includes the Explicit Extended CMTS MIC digest value, if present in the config file.

type, length, value for parameter 1
type, length, value for parameter 2
type, length, value for parameter n
type, length, value for Ext CMTS MIC Params
type, length, value for CM MIC

Figure D-4: Add CM MIC

- 4) Calculate the CMTS message integrity check (MIC) configuration setting as defined in clause D.2.1 and add to the file following the CM MIC using code and length values defined for this field and parameters defined in the Extended CMTS MIC Params configuration setting.

type, length, value for parameter 1
type, length, value for parameter 2
type, length, value for parameter n
type, length, value for Ext CMTS MIC Params
type, length, value for CM MIC
type, length, value for CMTS MIC

Figure D-5: Add CMTS MIC

- 5) Add the end of data marker and any needed padding bytes.

type, length, value for parameter 1	
type, length, value for parameter 2	
type, length, value for parameter n	
type, length, value for Ext CMTS MIC Params	
type, length, value for CM MIC	
type, length, value for CMTS MIC	
End of data marker	Pad (if required)

Figure D-6: Add End of Data Marker and Padding

D.1.3.1 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 TLVs which are not included in the CM MIC calculation:

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

These TLVs are the last TLVs in the CM configuration file.

NOTE: The bytes of the Extended CMTS MIC Params TLV are specifically included in the calculation and therefore needs to be inserted in the configuration file prior to the CM MIC. This includes the type, length and value fields.

The CM **MUST** accept configuration files with any number of TLVs following the CM MIC regardless of their length, unless the total file length exceeds the CM's maximum supported configuration file length.

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.2 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.2.1 CMTS MIC Calculation

The CMTS MUST calculate a CMTS MIC Digest value on TLVs of the REG-REQ/REG-REQ-MP message and compare it to the CMTS Message Integrity Check configuration setting in TLV7. If the Extended CMTS MIC Encoding is present but does not include an Explicit E-MIC Digest subtype, it indicates that the Extended CMTS MIC digest is implicitly provided in the CMTS MIC Configuration Setting of TLV7. In this case, the CMTS calculates only an Extended CMTS MIC digest using the TLVs indicated in the E-MIC Bitmap and compares it to the CMTS MIC Configuration Setting in TLV7. When the Extended CMTS MIC is implicitly provided in TLV7, the CMTS MUST confirm that the calculated Extended CMTS MIC digest matches the implicit digest in TLV7 in order to authorize the CM for registration.

If the Extended CMTS MIC Encoding is present and provides an Explicit E-MIC Digest subtype, the CMTS calculates both an Extended MIC Digest value and a "pre-3.0 DOCSIS" CMTS MIC digest value using the TLVs reported in REG-REQ or REG-REQ-MP. When both the Extended MIC digest and the pre-3.0 DOCSIS CMTS Digest are checked, the CMTS MUST consider a CM to be authorized when only the pre-3.0 DOCSIS CMTS Digest matches. If the pre-3.0 DOCSIS CMTS MIC digest matches but the explicit Extended CMTS MIC does not, the CMTS MUST silently ignore TLVs in REG-REQ and REG-REQ-MP which were marked as protected by the Extended CMTS MIC Bitmap and are not one of the pre-3.0 DOCSIS CMTS MIC TLVs provided in the Pre-3.0 DOCSIS CMTS MIC TLV List below.

If the Extended CMTS MIC Encoding TLV is not present or if the Extended CMTS MIC Encoding TLV is present and includes an Explicit E-MIC Digest Subtype, then the CMTS MUST calculate the message integrity check configuration setting by performing an MD5 digest over the following configuration setting fields when present in the REG-REQ or REG-REQ-MP messages, 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.
- DOCSIS Extension Field Configuration Settings (including Extended CMTS MIC Params).
- 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.
- Enable Test Modes.

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 extended CMTS MIC digest as stated in the clause D.2.1.1.

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 or REG-REQ-MP), regardless of its length.

On receipt of a configuration file containing an Extended CMTS MIC Encoding TLV, the CM MUST forward in the Registration Request message, all TLVs selected by the E-MIC Bitmap regardless of whether the CM understands the functionality related to those TLVs. The CM MUST send the TLVs that are selected for inclusion in the CMTS MIC or Extended CMTS MIC calculation in the order in which they appear in the config file.

It is important that the CM to preserve the ordering of TLVs from the config file, since this is the order in which they were used when calculating the Extended CMTS MIC Digest.

On receipt of a REG-REQ or REG-REQ-MP, the CMTS MUST validate the CMTS MIC. If the CMTS is unable to validate the REG-REQ or REG-REQ-MP according to the configuration setting (either because the REG-REQ or REG-REQ-MP does not contain the appropriate MIC TLV or because the HMAC type indicates a hash algorithm unsupported by the CMTS) the CMTS MUST reject the Registration Request by setting the authentication failure result in the Registration Response status field.

To validate the CMTS MIC, 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.

The CMTS MUST silently ignore any configuration file TLV in the Registration Request that is neither MIC protected (via the Pre-3.0 DOCSIS CMTS MIC or Extended CMTS MIC) nor one of the allowed unprotected TLVs explicitly mentioned in the list of "Configuration File Settings" provided in clause 6.4.7. As a result of this requirement, the configuration file generator needs to ensure that any configuration file TLV (other than those explicitly listed in clause 6.4.7) that is intended to be transmitted to the CMTS in Registration is protected by either the Pre-3.0 DOCSIS CMTS MIC or the Extended CMTS MIC or both.

D.2.1.1 Pre-3.0 DOCSIS CMTS MIC Digest Calculation

If the Extended CMTS MIC Configuration Setting TLV is not present or the Extended CMTS MIC Encoding is present and contains an Explicit Extended CMTS MIC Subtype, then the CMTS calculates a pre-3.0 DOCSIS CMTS MIC digest field using HMAC-MD5 as defined in [36] and only the set of pre-3.0 DOCSIS CMTS MIC TLVs in the order specified in clause D.2.1. When the CMTS calculates a pre-3.0 DOCSIS CMTS MIC digest, the CMTS MUST consider a CM to be unauthorized to register when its calculated pre-3.0 DOCSIS CMTS MIC Digest value differs from the CMTS MIC Configuration Setting in TLV 7 of a REG-REQ or REG-REQ-MP message.

D.2.1.2 Extended CMTS MIC Digest Calculation

When the Extended CMTS MIC Encoding is present, the CMTS MUST calculate the Extended CMTS MIC over the set of TLVs in REG-REQ or REG-REQ-MP as indicated by the Extended CMTS MIC Bitmap subtype. The CMTS MUST calculate the Extended CMTS MIC digest over the selected TLVs in the order that they were received in the Registration Request. Within Type fields, the CMTS MUST calculate the extended CMTS MIC digest 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.

If the Extended CMTS MIC Encoding is present in the REG-REQ/REG-REQ-MP message and no Explicit EMIC Digest subtype is provided, the CMTS MIC Configuration Setting in TLV7 is considered to "implicitly" provide an Extended CMTS MIC digest value. With an implicitly provided Extended CMTS MIC digest, the CMTS MUST compare the TLV7 CMTS MIC digest value to the calculated Extended CMTS MIC digest value. With implicit Extended CMTS MIC comparison, the CMTS MUST consider the CM to be unauthorized if the Extended CMTS MIC digest comparison fails.

The CMTS MUST support a configuration for the shared secret for Extended CMTS MIC calculation to differ from the shared secret for pre-3.0 DOCSIS CMTS MIC calculation, which uses the relatively insecure MD5 algorithm. In the absence of such configuration, the CMTS MUST use the same shared secret for Extended CMTS MIC Digest calculation as for pre-3.0 DOCSIS CMTS MIC digest calculation. The CMTS MUST calculate the Extended CMTS MIC using the algorithm specified in the Extended CMTS MIC Algorithm subtype. The CMTS MUST support the use of both the HMAC- MMH16- σ -n and the HMAC-MD5 hashing algorithms (see [15] for details of the MMH hash). The CMTS MAY support other hashing algorithms.

MMH is the preferred algorithm for DOCSISv3.0 (see [15]).

If the Explicit Extended CMTS MIC Digest Subtype is present, the CMTS compares its calculated E-MIC value to the Explicit E-MIC Digest value. If the Explicit Extended CMTS MIC Digest Subtype is present and the Extended CMTS MIC Bitmap indicates that TLV 43 is covered by the Extended CMTS MIC, the CMTS MUST copy the Extended CMTS MIC Digest value out of the Explicit Extended CMTS MIC Digest Subtype (TLV 43.6.3) and replace its value with zeros (0) prior to calculating the E-MIC value.

If the CMTS is unable to verify the Extended CMTS MIC digest, it MUST ignore TLVs in REG-REQ and REG-REQ-MP that are protected only by the Extended CMTS MIC.

Annex E (normative): Standard Receive Channel Profile Encodings

The following tables depict the verbose encodings of the standard receive channel profiles. Cable modems that support the 6 MHz RCP Center Frequency Spacing MUST support the profile with the RCP Name "CLAB-6M-004". Cable modems that support the 6 MHz RCP Center Frequency Spacing and also advertise a Multiple Receive Channel Support capability of 8 or greater (as defined in clause C.1.3.1.29) MUST also support the profile "CLAB-6M-008". Cable modems that support the 8 MHz RCP Center Frequency Spacing MUST support the profile with the RCP Name "CLAB-8M-004". Cable modems that support the 8 MHz RCP Center Frequency Spacing and also advertise a Multiple Receive Channel Support capability of 8 or greater (as defined in clause C.1.3.1.29) MUST also support the profile "CLAB-8M-008".

Table E-1: 2 Channel Standard Receive Channel Profile for 6 MHz DOCSIS (108 MHz to 870 MHz)

Type	Length	Value	Name
48	50		Receive Channel Profile
48.1	5	0x0010000002	Receive Channel Profile ID
48.2	11	"CLAB-6M-002"	RCP Name
48.3	1	6	RCP Center Frequency Spacing
48.4	6		Receive Module 1
48.4.1	1	1	Receive Module Index
48.4.2	1	10	Receive Module Adjacent Channels
48.5	9		Receive Channel
48.5.1	1	1	RC Index
48.5.2	1	0x40	RC Connectivity
48.5.5	1	1	RC Primary Downstream Channel Capable
48.5	6		Receive Channel
48.5.1	1	2	RC Index
48.5.2	1	0x40	RC Connectivity

Table E-2: 3 Channel Standard Receive Channel Profile for 6 MHz DOCSIS (108 MHz to 870 MHz)

Type	Length	Value	Name
48	58		Receive Channel Profile
48.1	5	0x0010000003	Receive Channel Profile ID
48.2	11	"CLAB-6M-003"	RCP Name
48.3	1	6	RCP Center Frequency Spacing
48.4	6		Receive Module 1
48.4.1	1	1	Receive Module Index
48.4.2	1	10	Receive Module Adjacent Channels
48.5	9		Receive Channel
48.5.1	1	1	RC Index
48.5.2	1	0x40	RC Connectivity
48.5.5	1	1	RC Primary Downstream Channel Capable
48.5	6		Receive Channel
48.5.1	1	2	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	3	RC Index
48.5.2	1	0x40	RC Connectivity

Table E-3: 4 Channel Standard Receive Channel Profile for 6 MHz DOCSIS (108 MHz to 870 MHz)

Type	Length	Value	Name
48	66		Receive Channel Profile
48.1	5	0x0010000004	Receive Channel Profile ID
48.2	11	"CLAB-6M-004"	RCP Name
48.3	1	6	RCP Center Frequency Spacing
48.4	6		Receive Module 1
48.4.1	1	1	Receive Module Index
48.4.2	1	10	Receive Module Adjacent Channels
48.5	9		Receive Channel
48.5.1	1	1	RC Index
48.5.2	1	0x40	RC Connectivity
48.5.5	1	1	RC Primary Downstream Channel Capable
48.5	6		Receive Channel
48.5.1	1	2	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	3	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	4	RC Index
48.5.2	1	0x40	RC Connectivity

Table E-4: 4 Channel Standard Receive Channel Profile for 6 MHz DOCSIS (108 MHz to 1 002 MHz)

Type	Length	Value	Name
48	80		Receive Channel Profile
48.1	5	0x0010000005	Receive Channel Profile ID
48.2	11	"CLAB-6M-005"	RCP Name
48.3	1	6	RCP Center Frequency Spacing
48.4	20		Receive Module 1
48.4.1	1	1	Receive Module Index
48.4.2	1	10	Receive Module Adjacent Channels
48.4.3	12		Receive Module Channel Block Range
48.4.3.1	4	111000000	Receive Module Minimum Center Frequency
48.4.3.2	4	999000000	Receive Module Maximum Center Frequency
48.5	9		Receive Channel
48.5.1	1	1	RC Index
48.5.2	1	0x40	RC Connectivity
48.5.5	1	1	RC Primary Downstream Channel Capable
48.5	6		Receive Channel
48.5.1	1	2	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	3	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	4	RC Index
48.5.2	1	0x40	RC Connectivity

Table E-5: 2 Channel Standard Receive Channel Profile for 8 MHz DOCSIS (108 MHz to 862 MHz)

Type	Length	Value	Name
48	50		Receive Channel Profile
48.1	5	0x0010001002	Receive Channel Profile ID
48.2	11	"CLAB-8M-002"	RCP Name
48.3	1	8	RCP Center Frequency Spacing
48.4	6		Receive Module 1
48.4.1	1	1	Receive Module Index
48.4.2	1	8	Receive Module Adjacent Channels
48.5	9		Receive Channel
48.5.1	1	1	RC Index
48.5.2	1	0x40	RC Connectivity
48.5.5	1	1	RC Primary Downstream Channel Capable
48.5	6		Receive Channel
48.5.1	1	2	RC Index
48.5.2	1	0x40	RC Connectivity

Table E-6: 3 Channel Standard Receive Channel Profile for 8 MHz DOCSIS (108 MHz to 862 MHz)

Type	Length	Value	Name
48	58		Receive Channel Profile
48.1	5	0x0010001003	Receive Channel Profile ID
48.2	11	"CLAB-8M-003"	RCP Name
48.3	1	8	RCP Center Frequency Spacing
48.4	6		Receive Module 1
48.4.1	1	1	Receive Module Index
48.4.2	1	8	Receive Module Adjacent Channels
48.5	9		Receive Channel
48.5.1	1	1	RC Index
48.5.2	1	0x40	RC Connectivity
48.5.5	1	1	RC Primary Downstream Channel Capable
48.5	6		Receive Channel
48.5.1	1	2	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	3	RC Index
48.5.2	1	0x40	RC Connectivity

Table E-7: 4 Channel Standard Receive Channel Profile for 8 MHz DOCSIS (108 MHz to 862 MHz)

Type	Length	Value	Name
48	66		Receive Channel Profile
48.1	5	0x0010001004	Receive Channel Profile ID
48.2	11	"CLAB-8M-004"	RCP Name
48.3	1	8	RCP Center Frequency Spacing
48.4	6		Receive Module 1
48.4.1	1	1	Receive Module Index
48.4.2	1	8	Receive Module Adjacent Channels
48.5	9		Receive Channel
48.5.1	1	1	RC Index
48.5.2	1	0x40	RC Connectivity
48.5.5	1	1	RC Primary Downstream Channel Capable
48.5	6		Receive Channel
48.5.1	1	2	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	3	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	4	RC Index
48.5.2	1	0x40	RC Connectivity

Table E-8: 4 Channel Standard Receive Channel Profile for 8 MHz DOCSIS (108 MHz to 1 002 MHz)

Type	Length	Value	Name
48	80		Receive Channel Profile
48.1	5	0x0010001005	Receive Channel Profile ID
48.2	11	"CLAB-8M-005"	RCP Name
48.3	1	8	RCP Center Frequency Spacing
48.4	20		Receive Module 1
48.4.1	1	1	Receive Module Index
48.4.2	1	8	Receive Module Adjacent Channels
48.4.3	12		Receive Module channel Block range
48.4.3.1	4	112000000	Receive Module Minimum Center Frequency
48.4.3.2	4	998000000	Receive Module Maximum center Frequency
48.5	9		Receive Channel
48.5.1	1	1	RC Index
48.5.2	1	0x40	RC Connectivity
48.5.5	1	1	RC Primary Downstream Channel Capable
48.5	6		Receive Channel
48.5.1	1	2	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	3	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	4	RC Index
48.5.2	1	0x40	RC Connectivity

Table E-9: 8 Channel Standard Receive Channel Profile for 6 MHz DOCSIS (108 MHz to 1 002 MHz)

Type	Length	Value	Name
48	112		Receive Channel Profile
48.1	5	0x0010000008	Receive Channel Profile ID
48.2	11	"CLAB-6M-008"	RCP Name
48.3	1	6	RCP Center Frequency Spacing
48.4	20		Receive Module 1
48.4.1	1	1	Receive Module Index
48.4.2	1	10	Receive Module Adjacent Channels
48.4.3	12		Receive Module channel Block range
48.4.3.1	4	111000000	Receive Module Minimum Center Frequency
48.4.3.2	4	999000000	Receive Module Maximum center Frequency
48.5	9		Receive Channel
48.5.1	1	1	RC Index
48.5.2	1	0x40	RC Connectivity
48.5.5	1	1	RC Primary Downstream Channel Capable
48.5	6		Receive Channel
48.5.1	1	2	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	3	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	4	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	5	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	6	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	7	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	8	RC Index
48.5.2	1	0x40	RC Connectivity

Table E-10: 8 Channel Standard Receive Channel Profile for 8 MHz DOCSIS (108 MHz to 1 002 MHz)

Type	Length	Value	Name
48	112		Receive Channel Profile
48.1	5	0x0010001008	Receive Channel Profile ID
48.2	11	"CLAB-8M-008"	RCP Name
48.3	1	8	RCP Center Frequency Spacing
48.4	20		Receive Module 1
48.4.1	1	1	Receive Module Index
48.4.2	1	8	Receive Module Adjacent Channels
48.4.3	12		Receive Module channel Block range
48.4.3.1	4	112000000	Receive Module Minimum Center Frequency
48.4.3.2	4	998000000	Receive Module Maximum center Frequency
48.5	9		Receive Channel
48.5.1	1	1	RC Index
48.5.2	1	0x40	RC Connectivity
48.5.5	1	1	RC Primary Downstream Channel Capable
48.5	6		Receive Channel
48.5.1	1	2	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	3	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	4	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	5	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	6	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	7	RC Index
48.5.2	1	0x40	RC Connectivity
48.5	6		Receive Channel
48.5.1	1	8	RC Index
48.5.2	1	0x40	RC Connectivity

Annex F (normative): The DOCSIS MAC/PHY Interface (DMPI)

F.1 Scope

Integrated Circuit (IC) chip sets with separate MAC and PHY chips used in the implementation of a CMTS SHOULD implement DMPI. DMPI does not apply to IC chip sets which integrate MAC and PHY components together into one chip.

Any usage of "MUST", "SHOULD" or "MAY" within the DMPI specification applies only if DMPI is implemented.

F.2 Conventions

F.2.1 Terminology

Throughout this annex, the terms MAC and PHY are used extensively. MAC is used to refer to the device which provides the interface between the PHY devices and the system. The term PHY refers to the device which performs the physical layer processing for a single RF channel. It is important to note that both of these terms refer to physical devices as opposed to layers in the IP protocol stack. For the purposes of this specification, integrated circuit chips which handle multiple RF channels simultaneously are considered to contain multiple PHY devices.

F.2.2 Ordering of Bits and Bytes

The following rules control the order of transmission of bits and bytes over all the interfaces specified in the document. In all cases, fields of Data Blocks are transmitted in the order in which they appear in the Data Block format description.

- Multi byte quantities are transmitted most significant byte first (big endian byte ordering). This byte ordering applies regardless of the width of the interface (byte, nibble, single bit).
- On nibble wide interfaces, the most significant nibble (bits 7:4) is transmitted first.
- On bit wide interfaces, the most significant bit of each field is transmitted first.

F.2.3 Signal Naming Conventions

Signal names which end with an "_N" are active low. Signals without this suffix are active high.

F.2.4 Active Clock Edge

All signals are driven and sampled on the rising edge of the clock except where otherwise noted.

F.2.5 Timing Specifications

The timing specs for DMPI use the following terminology.

Table F-1: Timing Parameters

Parameter	Symbol	Description
Clock Frequency	f	The frequency of the interface clock.
Clock Low Pulse Width	t_{lpw}	The low time of the interface clock.
Clock High Pulse Width	t_{hpw}	The high time of the interface clock.
Clock rise/fall time	t_{rf}	The transition time of the clock.
Input Setup Time to Clock	t_{su}	From when an interface signal is valid to the following rising clock edge.
Input Hold Time from Clock	t_h	From the rising clock edge to when an interface signal becomes invalid.
Clock to Signal Valid Delay	t_{cq}	From the rising edge of the interface clock to an interface signal becoming valid.

Following are some usage notes for these timing parameters:

- Setup and hold time specifications are given from the point of view of the DMPI Interface and not from the point of view of a device on the DMPI Interface. The clock to output, on the other hand, specifies the timing requirement of a DMPI device.
- The t_{su} parameter specifies the minimum guaranteed amount of setup time provided by the DMPI interface measured at the receiving device. Therefore, inputs on DMPI devices should require no more than this amount of setup time.
- The t_h parameter specifies the minimum guaranteed amount of hold time provided by the DMPI interface measured at the receiving device. Therefore, inputs on DMPI devices should require no more than this amount of hold time.
- The t_{cq} parameter specifies the minimum and maximum clock to output time at the driving device. The purpose of the minimum specification is to allow for clock skew between the driving and receiving DMPI device. For example, a 1ns minimum spec and a 0ns DMPI hold time requirement allows for at most 1ns of clock skew between devices. The maximum specification is to allow for the settling time of signals from the driving device to the receiving device and clock skew between devices.

F.3 Overview

This annex describes the DOCSIS MAC/PHY Interface (DMPI). DMPI is used to connect a DOCSIS MAC device to DOCSIS downstream and upstream PHY devices. While DMPI is a single interface, for the purposes of clarity, DMPI signals have been grouped into four separate groups. Each group serves a specific purpose and is independent of the others. For this reason, each group of signals is also referred to as an interface.

A Downstream PHY MUST include a Downstream Data Interface and an SPI Bus Interface. An Upstream PHY must include an Upstream Data Interface, an Upstream Control Interface and an SPI Bus Interface. PHY Chips which integrate multiple PHYs into a single package MUST have one set of interfaces for each PHY which has been integrated with the following exception.

An integrated PHY device MAY use a single select and a single SPI Bus for all internal PHYs (using the SPI Bus protocol described in clause F.9.4). An integrated Upstream PHY device MAY have only one TS_CLK input and only one US_CLK input.

A MAC MUST include one Downstream Data Interface for each Downstream PHY it supports and one set of Upstream Interfaces (Upstream Data and Upstream Control) for each Upstream PHY it supports. It MUST include at least one SPI Bus Interface.

DMPI has been defined with the following goals in mind:

- Vendor independence.
- Flexibility for future growth and vendor differentiation.

- Minimization of PHY specific logic in the MAC.

Figure F-1 shows an example application of DMPI.

NOTE: This figure shows the connections required for a single DS PHY and a single US PHY. Obviously, other applications with multiple DS and US PHYs are possible.

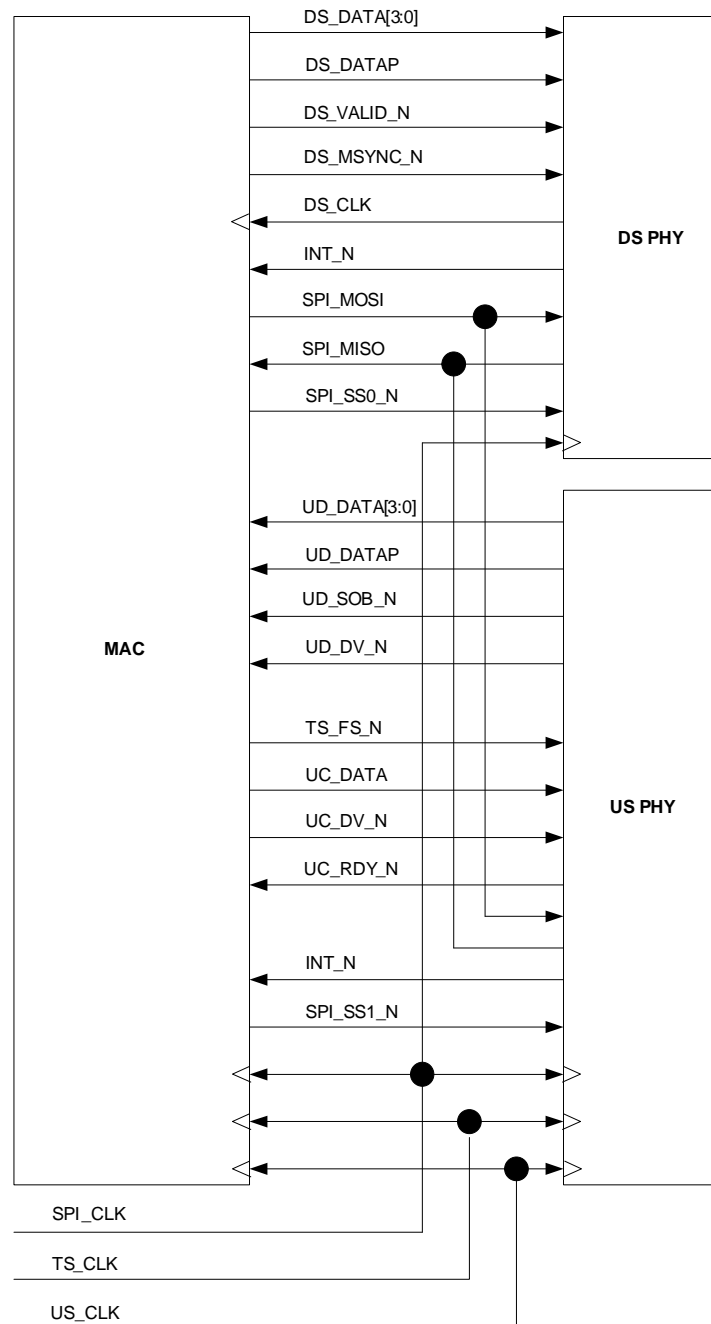


Figure F-1: DMPI Application

F.3.1 Downstream Data

The Downstream Data Interface carries data from the MAC to the PHY for transmission on the Downstream. All signals on the interface are synchronous with respect to a clock driven by the PHY and received by the MAC. Four bits of data are transferred on each clock. The frequency of this clock is proportional to the Downstream bit rate. Its precise frequency is a function of the Downstream Symbol Rate, the modulation type (64 QAM or 256 QAM) and the physical layer framing in use (ITU-T Recommendation J.83, annex A [i.2] or ITU-T Recommendation J.83, annex B [i.3]).

F.3.2 Upstream Data

The Upstream Data Interface carries data from the PHY to the MAC which has been received on the Upstream. The interface is synchronous to a dedicated interface clock whose frequency is not directly related to the upstream bit rate.

Data is transferred over the interface using a mixture of TLVs and TVs (a TLV for which the length is implied by the type). Along with the DOCSIS burst data, certain status information about the burst is also transferred to the MAC. There is also a TLV which allows the PHY to indicate that it did not receive a burst when one was expected.

F.3.3 Upstream Control

The Upstream Control Interface is used for two purposes. The first is to initialize the PHY's timestamp counter, frame counter and mini-slot counter and to check that the PHY's timestamp counter remains synchronized to the MAC's during operation. The second is to allow the MAC to pass information to the PHY regarding upcoming bursts.

This interface uses two clocks. The clock used for the counter synchronization is the 10,24 MHz CMTS master clock. A single signal that is synchronized to this clock is used to perform this counter synchronization. The other clock used for this interface is shared with the Upstream Data Interface and has a frequency unrelated to the upstream modulation clock or the 10,24 MHz CMTS master clock. This clock, along with an associated set of signals, is used to transfer descriptions of future bursts.

F.3.4 SPI Bus

The Serial Peripheral Interconnect (SPI) Bus is used to read and write registers in the PHYs. The system MAY use one or more SPI Buses to provide register access to the PHYs. The number of SPI Buses in the system is a function of the system's SPI Bus performance requirements. Each SPI Bus has a single master device which MAY be the MAC. Alternatively, an SPI Bus master MAY be some other device in the system (e.g. a microprocessor). References to the SPI Bus in this specification assume that the MAC is the master. The PHYs MUST only be slave devices. Each PHY MUST have one SPI Bus Interface. Multiple PHYs MAY share the same SPI Bus.

The SPI Bus definition includes an interrupt signal (INT_N). Each PHY MUST drive an interrupt. The interrupt signals MAY be received by an SPI Bus master or by some other device in the system which provides the ability to monitor their state.

F.4 Signals

F.4.1 Downstream Data

The signals used for the Downstream Data Interface are defined in table F-2.

Table F-2: Downstream Data Interface Signals

Signal	Description
DS_CLK	DS transmit clock Driven by the Downstream PHY See clauses F.6.1 and F.8.1 for detailed requirements for this clock
DS_MSNC_N	Downstream MPEG Sync Driven by the MAC marks first nibble of sync byte; active low
DS_VALID_N	Downstream Data Valid Driven by the MAC indicates that valid data is present on DS_DATA
DS_DATA[3:0]	DS transmit data Driven by the MAC
DS_DATAP	Downstream Parity Driven by the MAC Even parity for DS_DATA (the number of 1's across DS_DATA and DS_DATAP is even) DS_DATA and its corresponding DS_DATAP are driven on the same clock. Parity is not delayed a clock as it is in some interfaces

F.5 Upstream Data

The signals used for the Upstream Data Interface are defined in table F-3.

Table F-3: Upstream Data Interface Signals

Signal	Description
US_CLK	Upstream Data/Control Clock Driven by external clock source (input to MAC and PHY)
UD_SOB_N	Upstream Data Start of Data Block Driven by Upstream PHY asserted when the first nibble or first byte of the Data Block is on UD_DATA
UD_DV_N	Upstream Data Valid Driven by Upstream PHY indicates valid data on UD_DATA
UD_DATA[3:0]	Upstream Data Driven by Upstream PHY
UD_DATAP	Upstream Data Parity Driven by Upstream PHY Even parity for UD_DATA (the number of 1's across UD_DATA and UD_DATAP is even) UD_DATA and its corresponding UD_DATAP are driven on the same clock. Parity is not delayed a clock as it is in some other interfaces

F.5.1 Upstream Control

Table F-4 lists the signals that are used for the Upstream Control Interface.

Table F-4: Upstream Control Interface Signals

Signal	Description
US_CLK	Upstream Clock Driven by external clock source (input to MAC and PHY)
UC_DV_N	Upstream Control Data Valid Driven by the MAC Indicates valid Upstream Control Message data on UC_DATA
UC_DATA	Upstream Control Data Driven by the MAC
UC_RDY_N	Upstream Control Ready Driven by the PHY Indicates that the PHY is ready to receive an Interval Description Message
TS_CLK	10,24 MHz master clock Driven by external clock source (input to MAC and PHY)
TS_FS_N	Timestamp Frame Sync Driven by the MAC

F.5.2 SPI Bus

Table F-5: SPI Bus Signals

Signal name	Description
SPI_CLK	SPI Bus Clock Driven by a source external to the MAC and PHY or driven by the MAC
SPI_MOSI	Master out/Slave in Serial data from the MAC to the PHY
SPI_MISO	Slave out/Master in Serial data from the PHY to the MAC MAY be driven by the PHY from the falling edge of SPI_CLK
SPI_SSx_N	Slave Select Selects a slave for a transaction One Slave Select signal is provided by the MAC for each PHY (x = 1 to N). Addressing of devices within a package is provided by the protocol layer described in clause F.9.4. MAY be sampled by the PHY on the falling edge of SPI_CLK
INT_N	Interrupt Driven by PHYs Open drain

F.5.3 Parity

The Downstream Data, Upstream Data and Upstream Control Interfaces use parity to maintain data integrity on the interface. Parity SHOULD be implemented.

The SPI Bus does not have parity.

Parity is even and covers only the data lines of the interface. Specific rules for parity checking are detailed in the following clauses.

F.5.3.1 Downstream Data

Parity must be checked by the Downstream PHY and covers DS_DATA. Since the Downstream transmit data is protected (DOCSIS frame HCS and CRC), detection of a parity error is not considered fatal and MUST NOT cause the processing of transmit data to halt. The PHY must generate an interrupt to the system when it detects a parity error so that the system can be made aware of its occurrence. Parity checking on this interface provides a way to distinguish between data errors on the interface and those in other parts of the data path.

Parity is checked by the MAC and covers UD_DATA. Since the Upstream receive data is protected (DOCSIS frame HCS and CRC), detection of a parity error is not considered fatal and **MUST NOT** cause the processing of receive data to halt. The MAC must generate an interrupt to the system when it detects a parity error so that the system can be made aware of its occurrence. Parity checking on this interface provides a way to distinguish between data errors on the interface and those in other parts of the data path.

Parity is checked by the Upstream PHY and covers the entire Upstream Control Message. A parity error on this interface is considered a fatal error. The PHY MUST NOT process the Upstream Control message which was received with a parity error as well as any subsequently received message. The PHY MAY process any Upstream Control messages received prior to the occurrence of the parity error. This processing MAY include the passage of various types of Upstream Data Blocks to the MAC.

- One active low interrupt line of level type.
- Driven open drain.
- Cause of interrupt line assertion determined by software read(s) of PHY register(s) which contains one bit for each interrupt source.
- No hardware prioritization of interrupt sources.
- Each interrupt source separately cleared by software write(s) to PHY register(s).
- Asserted until all interrupt source bits are cleared (interrupt line is a simple OR of all interrupt sources).

DS_CLK

DS_MS SYNC_N

DS_VALID_N

DS_DATA

DS_DATAP

ETSI

The following behavior of DS_CLK and DS_VALID_N is required:

- DS_CLK MUST NOT be gapped (it must have a constant frequency).
- DS_CLK frequency MUST be 1/4 of the Downstream Line Rate. The DS Line Rate is the data rate including the ITU-T Recommendation J.83 annex A [i.2] framing overhead.
- The MAC MUST assert DS_VALID_N for the entire 188 byte MPEG packet transfer and then MUST de-assert it for exactly 32 clocks following the transfer of the last nibble of the MPEG packet.

F.6.2 Downstream Data (ITU-T Recommendation J.83, annex B)

Figure F-3 shows the protocol used to transfer data across this interface for ITU-T Recommendation J.83, annex B [i.3] operation.




Figure F-3: Downstream Data Signal Protocol for annex B Operation

The following behavior of DS_CLK and DS_VALID_N is required:

- DS_CLK MUST NOT be gapped (it must have a constant frequency).
- DS_CLK frequency MUST be 1/4 of the Downstream Payload Rate. The Downstream Payload Rate is the data rate excluding the ITU-T Recommendation J.83, annex B [i.3] framing overhead.
- The MAC MUST keep DS_VALID_N always asserted.

F.6.3 Upstream Data

Figure F-4 shows the signalling protocol for this interface.




Figure F-4: Upstream Data Protocol

It is a very simple protocol in which the Upstream PHY indicates the presence of valid data on UD_DATA by asserting UD_DV_N. The MAC has no ability to control the flow of data and is required to sample UD_DATA on every rising clock edge on which UD_DV_N is asserted. The start of a Data Block is indicated by the PHY's assertion of UD_SOB_N. This signal **MUST** be asserted when the first nibble of the first byte of the Data Block is driven onto UD_DATA.

The MAC **MUST** keep track of length of each Data Block as it relates to the assertion of UD_SOB_N. If UD_SOB_N is asserted before the entire previous Data Block has been transferred, the MAC **MUST** drop the associated burst and generate an interrupt.

If the FIRST_STATUS byte indicates the absence of a PHY_STATUS Data Block but the PHY transfers one, the PHY_STATUS Data Block **MUST** be discarded by the MAC and an error **MUST** be signaled to the system.

F.6.4 Upstream Control

F.6.4.1 Counter Synchronization

The master timestamp counter **MUST** reside in the MAC. The master mini-slot counter and master frame counter **MUST** reside in the PHY. The PHY **MUST** capture a timestamp snapshot on every frame boundary. When the system needs a Timestamp Snapshot for a UCD, it **MUST** read this snapshot using a single SPI bus transaction. The PHY **MUST** ensure that the timestamp snapshot does not change during the SPI Bus read transaction.

A common timestamp clock, TS_CLK, **MUST** be externally provided to the upstream PHYs and the MACs. The frequency of this timestamp clock **MUST** be 10,24 MHz \pm 5 ppm. The MAC **MUST** synchronize all PHYs to the timestamp value of the MAC. To accomplish this, the MAC **MUST** provide a frame sync pulse, TS_FS_N, to the PHYs that is synchronous to the positive edge of TS_CLK and has a pulse width equal to one period of TS_CLK.

The 32 bit timestamp counter consists of a group of upper bits and a group of lower bits. The MAC and PHY **MUST** provide at least the following choices of upper and lower bit boundaries shown in table F-6.

Table F-6: Timestamp Counter Initialization Options

Upper Bits	Lower Bits	Frame Sync Interval
8	24	1 638,4 ms
9	23	819,2 ms
10	22	409,6 ms
11	21	204,8 ms
12	20	102,4 ms

Figure F-5 shows an example of the proper assertion of the TS_FS_N signal.

NOTE: The **TIMESTAMP** is shown for reference and is not part of the Upstream Control Interface. In this example, Upper Bits = 8.




Figure F-5: Counter Synchronization

The MAC **MUST** assert TS_FS_N two 10,24 MHz clock periods prior to the lower bits of the MAC timestamp counter equaling all zeros. The MAC **SHOULD** provide some sort of maskable indication to the system when TS_FS_N occurs so that the system will have time to program the registers of the PHYs prior to the next assertion of TS_FS_N. The period of TS_FS_N is a function of the timestamp bit time and the number of lower bits from table F-6. The variation of the TS_FS_N period is to allow the system designer to trade off system response time versus the time available to initialize a PHY chip.

The PHY MUST provide all combinations of the following three initialization options when TS_FS_N is asserted:

- the upper bits of the timestamp counter are specified and the lower bits are set to zero;
- the full 8 bits of the frame counter are specified;
- the full 32 bits of the mini-slot counter are specified.

The specification of these counters is supplied across the SPI Bus prior to the next frame sync pulse. Two TS_CLK clock cycles after TS_FS_N occurs, the PHY chip MUST initialize the specified counters. These counters are loaded at configuration time and not on every assertion of TS_FS_N. A single PHY may be re-initialized without the need to re-initialize or otherwise interrupt the operation of other PHYs or the MAC.

During normal operation, the PHY MUST check that the lower bits of the PHY timestamp counter are exactly all zeros two 10,24 MHz clock cycles following every assertion of TS_FS_N. If the check is negative, the PHY MUST generate an interrupt and MUST provide status accessible over the SPI bus.

F.6.4.2 Upstream Control Messages

Figure F-6 shows a sample transaction.




Figure F-6: Upstream Control Message Transfer

The Upstream Control Interface is used to transfer time critical configuration information (messages) to the PHY. The most common type of message is an Interval Description message. This message informs the PHY of the arrival time and characteristics of an upcoming burst. The protocol of this interface is very simple. Following is a description of how this interface works:

- A transaction transfers a single Upstream Control message.
- UC_DV_N MUST remain asserted for the entire duration of the Upstream Control message transfer.
- The length of each Upstream Control message is inferred by its type.
- UC_DV_N MUST be de-asserted for a minimum of one US_CLK clock period to indicate the end of a transaction.
- UC_RDY_N MAY be used to stop and start the flow of Interval Description messages. UC_RDY_N does not affect the transfer of other message types. If the PHY is receiving an interval description and does not want to receive a subsequent interval description, the PHY MUST de-assert UC_RDY_N at least two clock cycles of US_CLK prior to the end of the current interval description. This de-assertion behavior is shown in figure F-6. The MAC MUST transfer a new Interval Description Message within 10 US_CLK periods of the assertion of UC_RDY_N if a new Interval Description Messages is available.

F.6.5 SPI Bus

Figure F-7 shows a SPI Bus transaction.




Figure F-7: SPI Bus Transaction

A transaction proceeds as follows:

- The master asserts the select (SPI_SSx_N) of the desired slave device.
- The master drives SPI_MOSI with the appropriate command and data as described in [1].
- For write commands, the first byte of data driven on SPI_MOSI is written to the register specified by the address in the command. The second byte of data (if it exists), is written to the next higher numbered address. Writes continue in this way until the master terminates the transaction by de-asserting SPI_SSx_N.
- For read commands, the slave drives the read data on SPI_MISO which is indicated by the address in the command. The first bit of this read data is driven one clock after the last bit of the command has been sampled. Read data from consecutively numbered addresses is driven until the master terminates the transaction by de-asserting SPI_SSx_N.

SPI_CLK MUST be driven (oscillate) for at least one clock period prior to the assertion of SPI_SSx_N, during the entire SPI Bus transaction and for one clock after the deassertion of SPI_SSx_N. SPI_CLK MAY be driven high or low at all other times.

F.7 Electrical Specifications

F.7.1 DC Specifications

Devices which connect to DMPI must meet the requirements listed in table F-7.

NOTE: Output High Voltage and Output High Current specifications do not apply to the INT_N output as it is open drain.

Table F-7: DC Characteristics

Parameter	Symbol	Min.	Max.	Units
Input Capacitance			10	pf
Input Low Voltage	Vil		0,8	v
Input High Voltage	Vih	2,0		v
Output Low Voltage	Vol		0,4	v
Output High Voltage	Voh	2,4		v
Output Low Current	Iol	4		ma
Output High Current	Ioh	-4		ma

F.8 Timing Specifications

F.8.1 Downstream Data

Table F-8: DS Data Interface Timing

Parameter	Symbol	Min.	Max.	Units
DS_CLK Frequency	f		25	MHz
DS_CLK Low Pulse Width	t_{lpw}	10		ns
DS_CLK High Pulse Width	t_{hpw}	10		ns
DS_CLK rise/fall time	t_{rf}		4	ns
DS_CLK Jitter	t_j		97,66	ns
Input Setup Time to DS_CLK	t_{su}	10		ns
Input Hold Time from DS_CLK	t_h	0		ns
DS_CLK to Signal Valid Delay	t_{cq}	1	15	ns

F.8.2 Upstream Data

Table F-9: US Data Interface Timing

Parameter	Symbol	Min.	Max.	Units
US_CLK Frequency	f	33	40,96	MHz
US_CLK Low Pulse Width	t_{lpw}	6,5		ns
US_CLK High Pulse Width	t_{hpw}	6,5		ns
US_CLK rise/fall time	t_{rf}		3	ns
Input Setup Time to US_CLK	t_{su}	6		ns
Input Hold Time from US_CLK	t_h	0		ns
US_CLK to Signal Valid Delay	t_{cq}	1	12	ns

F.8.3 Upstream Control

Table F-10: Upstream Control Interface Timing

Parameter	Symbol	Min.	Max.	Units
Input Setup Time to US_CLK	t_{su}	6		ns
Input Hold Time from US_CLK	t_h	0		ns
US_CLK to Signal Valid Delay	t_{cq}	1	12	ns
TS_CLK rise/fall time	t_{rf}		3	ns
Input Setup Time to TS_CLK	t_{su}	10		ns
Input Hold Time from TS_CLK	t_h	0		ns
TS_CLK to Signal Valid Delay	t_{cq}	1	15	ns

F.8.4 SPI Bus

Table F-11: SPI Bus Timing

Parameter	Symbol	Min.	Max.	Units
SPI_CLK Frequency	f		10,24	MHz
SPI_CLK Low Pulse Width	t _{lpw}	43,9		ns
SPI_CLK High Pulse Width	t _{hpw}	43,9		ns
SPI_CLK rise/fall time	t _{rf}		4	ns
SPI_MOSI or SPI_MISO Setup Time to SPI_CLK	t _{su}	15		ns
SPI_MOSI or SPI_MISO Hold Time from SPI_CLK	t _h	0		ns
SPI_SSx_N Setup Time to SPI_CLK rising	t _{su}	50		ns
SPI_SSx_N Setup Time to SPI_CLK falling	t _{su}	25		ns
SPI_SSx_N Hold Time from SPI_CLK	t _h	0		ns
SPI_CLK to Signal Valid Delay	t _{cq}	1	12	ns

F.9 Data Format and Usage

F.9.1 Downstream Data

The data which passes from the MAC to the PHY is a stream of MPEG packets. The start of the SYNC byte is indicated by the assertion of the DS_MSYNc_N signal. Including the SYNC byte, each MPEG packet is 188 bytes in length.

The MAC MUST generate null MPEG packets when there are no DOCSIS frames to be transmitted.

F.9.2 Upstream Data

F.9.2.1 Block Format

Data is passed from the Upstream PHY to the MAC using a combination variable sized units called Upstream Data Blocks. Each of these Data Blocks has the generic format described in table F-12, (except for the CHANNEL Data Block Type as indicated in clause F.9.2.8.5).

Table F-12: Upstream Data Block Format

Size (bytes)	Name	Description
1	Block Type	Identifies the type of Block
2	Block Length	Length of Block data field in bytes (N) Not present for CHANNEL Block Type
N	Block Data	Block data

As can be seen from this table, each Data Block starts with a Data Block type. This type is used by the MAC to determine which type of Data Block data is being transferred. The Data Block length field contains the length in bytes of the Data Block data and is used by the MAC to find the end of the Data Block data field. In most cases, the Data Block type determines the format of the Data Block data field. The exception to this is the PHY_STATUS type where the format of the Data Block data field is PHY specific.

Table F-13 gives a complete list of all Block Types.

Table F-13: Upstream Data Block Types

Type	Name	Description
0x00	Reserved	Reserved
0x01	FIRST_DATA	First data of burst Contains 7 bytes of fixed format status data and first data of burst
0x02	MIDDLE_DATA	Middle data of burst
0x03	LAST_DATA	Last data of burst Contains 4 bytes of fixed format status data and last data of burst
0x04	PHY_STATUS	Status which should be passed to software The maximum length of this Block is 128 bytes
0x05	NO_BURST	Indicates that no burst was received during a transmit opportunity
0x06	CHANNEL	Used to indicate the channel to which the next Data Block belongs
0x07 to 0xff	Reserved	Reserved

F.9.2.2 FIRST_DATA Block

Table F-14 shows the format of the FIRST_DATA Block.

The FIRST_DATA Block is used by the PHY to transfer the beginning of a received burst. This block MUST contain the seven bytes of status information defined in the table. It MAY contain burst data as well. The Block Length of the FIRST_DATA block MUST NOT be less than seven.

NOTE: N=7 is allowed.

Table F-14: FIRST_DATA Data Format

Size (bytes)	Name	Description
1	FIRST_STATUS	Bit 7:6, reserved, MUST be zero Bit 5, New UCD, 1 => First burst received on new UCD Bit 4, PHY_STATUS Data Block present, 1 => PHY_STATUS Data Block present Bit 3:0, IUC, taken from the Upstream Control Interval Description message
2	SID	Bit 15:14, reserved, MUST be zero Bit 13:0, SID, taken from the Upstream Control Interval Description message
4	START_MINISLOT	Derived from the Upstream Control Interval Description message parameters
N-7	BURST_DATA	First data of burst

F.9.2.3 MIDDLE_DATA Block

Table F-15 shows the format of the MIDDLE_DATA Block. The MIDDLE_DATA block is used to transfer burst data.

Table F-15: MIDDLE_DATA Data Format

Size (bytes)	Name	Description
N	BURST_DATA	Middle data of burst

F.9.2.4 LAST_DATA Block

Table F-16 shows the format of the LAST_DATA Block. The LAST_DATA block is used to transfer burst data. This block MUST contain the four bytes of status information defined in the table. It MAY also contain burst data. The Block Length of the LAST_DATA block MUST NOT be less than four.

NOTE: N=4 is allowed.

Table F-16: LAST_DATA Data Format

Size (bytes)	Name	Description
N-4	BURST_DATA	Last data of burst
1	LAST_STATUS	bit 7:3, reserved, must be zero bit 2, internal PHY error, 1 => internal PHY error bit 1, low energy; indicates that the burst power was below the desired threshold, 1 => low energy bit 0, high energy; indicates that the burst power was above the desired threshold, 1 => high energy
1	GOOD_FEC	the number of good FEC blocks in the burst must stop incrementing when count reaches 255 must be zero if FEC is disabled for associated interval
1	CORRECTED_FEC	the number of corrected FEC blocks in the burst must stop incrementing when count reaches 255 must be zero if FEC is disabled for associated interval
1	UNCORRECTED_FEC	the number of uncorrected FEC blocks in the burst must stop incrementing when count reaches 255 must be zero if FEC is disabled for associated interval

F.9.2.5 PHY_STATUS Block

Table F-17 shows the format of the PHY_STATUS Block. The PHY_STATUS block is used to transfer PHY unique status to the MAC. The contents of this block are vendor unique and are unrestricted.

Table F-17: PHY_STATUS Data Format

Size (bytes)	Name	Description
N	PHY_STATUS	PHY specific status information such as channel characteristics (e.g. timing error, power error, frequency error, EQ coefficients)

F.9.2.6 NO_BURST Block

Table F-18 shows the format of the NO_BURST Block. This block is used by the PHY to indicate that a valid burst was not received when one was expected. Absence of a valid burst may be caused by either no transmitter, multiple transmitters or a noise corrupted transmission. DMPI does not specify the criteria by which the PHY distinguishes between these cases.

Table F-18: NO_BURST Data Format

Size (bytes)	Name	Description
2	SID_STATUS	bit 15, collision, collision occurred bit 14, no energy, no energy detected bit 13:0, SID, taken from the Upstream Control Interval Description message
4	START_MINISLOT	Derived from the Upstream Control Interval Description message parameters
1	IUC	bit 7:5, reserved, must be zero bit 4: New UCD, 1 => First NO_BURST block received on new UCD bit 3:0, IUC, taken from the Upstream Control Interval Description message
2	LENGTH	Taken from the Upstream Control Interval Description message Note that for Contention Intervals, this is the length of the interval and not the length of each individual transmit opportunity in the interval

F.9.2.7 CHANNEL Block

Table F-19 shows the format of the CHANNEL Block. The Channel Block is used by the PHY to indicate to which logical channel subsequent blocks belong.

Table F-19: CHANNEL Data Format

Size (bytes)	Name	Description
1	CHANNEL	bit 7:3, reserved, must be zero bit 2:0, Channel Number

F.9.2.8 Block Usage

F.9.2.8.1 Overview

At least one Data Block **MUST** be transferred for every Transmit Opportunity. If a burst is received during a transmit opportunity, the appropriate series of Data Blocks **MUST** be transferred to the MAC (FIRST_DATA, MIDDLE_DATA, LAST_DATA, PHY_STATUS). If no burst is received, a NO_BURST Data Block **MUST** be transferred unless the region was allocated to a SID which the system has reserved for no CM (e.g. the null SID as defined in clause A.2.1).

NOTE: Since contention regions have multiple transmit opportunities, more than one set of Data Blocks will likely be transferred to over the interface for each region (interval).

The minimum amount of payload in a Data Block (the length of the Block Data field) **MUST** be 16 bytes with the following exceptions:

- Data Blocks for bursts which are less than 16 bytes in length.
- Any LAST_DATA Data Block.

The Upstream PHY **SHOULD** minimize the number of Data Blocks required to transfer a burst so as to minimize the amount of overhead on DMPI. However, nothing specific other than what is mentioned above is required.

For Non-contention Intervals, the START_MINISLOT **MUST** be equal to the START_MINISLOT which was passed to the PHY in the corresponding Interval Description Message (described in clause F.9.3.1). For Contention Intervals (IE types REQ and REQ/Data), the PHY **MUST** calculate an accurate START_MINISLOT value and return it in the appropriate Data Block (FIRST_DATA or NO_BURST). In general terms, this means that PHY **MUST** calculate the START_MINISLOT for each Data Block by taking into account the number of mini-slots which have passed since the start of the Interval. Specifically, the Upstream PHY **SHOULD** use the IUC and SID in the Upstream Control Interval Description message to calculate a burst start offset from the original START_MINISLOT value received in this message. The offset is then added to this START_MINISLOT and returned to the MAC as the START_MINISLOT in the appropriate Upstream Data Block.

F.9.2.8.2 Burst Data Transfer

The transfer of a burst **MUST** be accomplished by transferring the following Data Blocks in the following order:

- one FIRST_DATA Block;
- zero to N MIDDLE_DATA Blocks;
- one LAST_DATA Block;
- zero or one PHY_STATUS Block.

The only Data Block type which **MAY** be transferred after a FIRST_DATA Data Block and before a LAST_DATA Data Block is a MIDDLE_DATA Data Block. Any other Data Block transferred between these two Data Blocks **MUST** be discarded by the MAC.

In general, each Data Block will contain one FEC block of data. However, there is no specific requirement as to which Data Block types contain which parts of the burst data. The data MAY be distributed between the various Data Block types at the discretion of the PHY as long as the Data Block ordering shown above is maintained and the minimum block length requirements are respected. A Data Block type with a length of zero is also allowed. Every burst, regardless of size, MUST be transferred to the MAC using at least a FIRST_DATA Block and a LAST_DATA Data Block. The PHY_STATUS Data Block is optional with its presence indicated in the FIRST_STATUS byte in the FIRST_DATA Data Block. The MIDDLE_DATA Data Block is optional.

Typically, there will be some arbitrary delay between the transfer of one Data Block and the transfer of the next. It is the PHY's responsibility to assure that these delays do not interfere with the PHY's ability to keep up with the incoming data rate.

NOTE: This series of Data Blocks is passed to the MAC any time a burst is received regardless of the type of interval in which the burst was received (contention or non-contention).

F.9.2.8.3 No Burst Status Transfer

It is sometimes useful for the system to know when no usable burst was received during a transmit opportunity. This can happen when there is no transmitter (no energy) in the opportunity, there is more than one transmitter (a collision) or noise corrupted a transmission. For a contention region, knowledge of unused opportunities or those with collisions helps software optimize its scheduling of contention regions (their duration and frequency). For non-contention regions, these same events could be an indication of a problem with a CM. Or, they could be a result of illegal or malicious use of the US bandwidth.

The NO_BURST Data Block contains two status bits. The one called "collision" indicates that a collision occurred during the transmit opportunity. The other, called "no energy", indicates that there was no energy detected during the transmit opportunity. If neither is set, it means that there was energy but that no preamble was found. Both of these bits must not be set at the same time.

F.9.2.8.4 UCD Change Indication

In order to allow the system to properly size grants for bandwidth requests which were received prior to a UCD change but are granted after such a UCD change, the MAC needs to be notified that a new UCD is in effect. This notification is achieved via "New UCD" status bits in the NO_BURST and FIRST_DATA Data Blocks. The PHY MUST set the New UCD bit of the first Data Block sent to the MAC after a UCD change (FIRST_DATA or NO_BURST, whichever is sent first). The New UCD bit of these Data Blocks MUST be zero at all other times.

F.9.2.8.5 Logical Channel Support

For Upstream PHYs which support multiple logical channels, a Data Block Type called CHANNEL is used to specify to which logical channel each Data Block belongs. This Data Block contains a single byte of payload which is the channel number (zero to seven inclusive). Since the Data Block is a fixed length and is potentially required for every other Data Block transferred, the length bytes are omitted from the normal Data Block format and only the Data Block Type and Block Data are transferred. So, a CHANNEL Block is always two bytes long (including the Type byte).

It is important to note that the Channel Data Block is only used to distinguish between data received on logical channels within the same RF channel. Since each PHY has its own DMPI interface, the RF channel to which data belongs is inferred by the PHY's connection to the MAC.

The CHANNEL Data Block is used as follows:

- The CHANNEL Data Block sets the "current" channel for transmitted Data Blocks. After reset, the MAC must set the current channel to zero.
- The current channel is always the channel number contained in the most recently transmitted CHANNEL Data Block. For this reason, transmission of a CHANNEL Data Blocks is only required when a change in the current channel is desired.

Since the MAC sets the current channel to zero prior to receipt of any CHANNEL Data Blocks, PHYs which support a single channel are not required to support this Data Block Type. In cases where multiple CHANNEL Data Blocks are transferred in succession, the last one received prior to the transfer of one of the other Data Blocks will be considered valid and the others that preceded it will be ignored. NO_BURST Data Blocks may be preceded by a CHANNEL Data Block. If a series of NO_BURST Data Blocks for the same channel are transmitted to the MAC, only one CHANNEL Data Block is required (transferred prior to the first NO_BURST Data Block).

All Data Blocks associated with a single burst MUST be transferred contiguously over the Upstream Data interface. Specifically, this would mean that FIRST_DATA, MIDDLE_DATA, LAST_DATA, PHY_STATUS would all be transferred for a given burst of a given channel before any other Data Blocks were transferred for another channel. A CHANNEL Data Block MUST precede the first Data Block (NO_BURST or FIRST_DATA) that belongs to a channel which is different than the one which preceded it. The PHY MAY transfer a CHANNEL Data Block prior to the FIRST_DATA block of every burst. CHANNEL Data Blocks MUST NOT be transferred immediately before any of the other Data Block associated with a burst.

F.9.3 Upstream Control

The Upstream Control Interface carries two different messages. One of them is used to describe upcoming bursts. The other is used to indicate UCD changes.

The format of an Upstream Control Message is shown in table F-20.

Table F-20: Upstream Control Message Format

Size (bits)	Name	Description
3	TYPE	Message Type
3	CHANNEL	Logical Channel Number
N	PAYLOAD	Payload of Message
1	PARITY	Even parity for all bits in the Message (the number of 1's across all bits in {TYPE, CHANNEL, PAYLOAD, PARITY} is even)

Table F-21 shows the Message Type encoding.

Table F-21: Upstream Message Types

Type	Name	Description
0x0	INTERVAL_DESCRIPTION	Describes an interval
0x1	UCD_CHANGE	Indicates a UCD change has occurred
0x2 to 0x7	Reserved	Reserved

F.9.3.1 Interval Description Message

Table F-22 describes the format of the Interval Description Message Payload.

Table F-22: Upstream Interval Description Format

Size (bits)	Name	Description
14	SID	expected SID from MAP IE
4	IUC	IUC from MAP IE
14	LENGTH	length in mini-slots
32	START_MINISLOT	starting mini-slot of interval (alloc start time + offset of IE)
3	PSC	PHY_STATUS Control

The MAC builds these Interval Description Messages from the information present in the DOCSIS MAPs that have been generated for the logical channels which the PHY is servicing. The MAC MUST transfer only one Interval Description Message to the PHY for an Interval Allocation which might describe an interval which has more than one transmit opportunity (e.g. REQ, REQ/DATA). The MAC MAY generate Interval Description Messages for Interval Allocations to the NULL SID. The MAC MUST NOT generate Interval Description Messages for Interval Allocations for the NULL SID if they overlap with Interval Allocations for non-NULL SIDs on other logical channels being serviced by the same PHY. Interval Description Messages for the NULL SID MAY be associated with any currently active logical channel. In contrast to MAP messages, the set of all Interval Description Messages from the MAC taken together need not describe every mini-slot on the logical channels in question; the MAC MAY refrain from sending Interval Description messages to describe inactive time periods on any or all logical channels. So as to minimize the complexity and buffering requirements of the PHY, the MAC MUST sort the Interval Descriptions from all logical channels, putting them into chronological order and deliver them to the PHY in this order. Note that Interval Description Messages MUST NOT be transferred for the NULL IE, Data Acknowledgement IE's or Data Grants Pending (since none of these is an Interval Allocation).

The system is allowed to schedule the Initial Maintenance regions of all logical channels of a physical channel to occur simultaneously. This type of overlap MUST be handled as follows:

- The MAC MUST transfer an Interval Description for only one of the logical channels.
- The Interval Description which is transferred MUST be the one with the earliest start time. If more than one Interval Description has the earliest start time, the MAC MAY choose any of these overlapping interval descriptions to pass to the PHY.
- The PHY MUST accept any of the logical channel numbers it supports for this interval description.

The system software is responsible for knowing that bursts received during initial ranging could be from CMs on any of the logical channels.

It is possible for there to be illegal overlap of intervals for the logical channels. An illegal overlap is defined to be an overlap of intervals other than Initial Maintenance. The PHY MAY detect these illegal overlaps. If the PHY performs this function, it MUST generate an interrupt to alert the system of such an event. It MUST capture the illegally overlapped Interval Description and hold it in SPI Bus accessible registers until software acknowledges its receipt.

The PSC field of the Interval Description Message is used to control the contents of the PHY_STATUS block. The usage of this field is summarized below:

- If PSC = 000, the contents of the PHY_STATUS block is determined through PHY programmable registers.
- If PSC is any other value, the contents of the PHY_STATUS block are vendor specific.

The MAC and PHY MUST support PSC = 000.

The MAC and PHY MAY support other values.

F.9.3.2 UCD Change Message

Table F-23 describes the format of the UCD Change Message Payload.

Table F-23: UCD Change PAYLOAD Format

Size (bits)	Name	Description
8	CCC	Configuration Change Count from the MAP

The MAC MUST send this message before sending the first Interval Description message after a UCD change. This message MUST NOT be sent at any other time.

F.9.4 SPI Bus

In order to perform an SPI Bus transaction, the master MUST drive SPI_MOSI with a bitstream of the following format.

Table F-24: SPI Bus Transaction Format

Size (bits)	Name	Description
4	DEVICE_ID	Device ID
3	RSVD	Reserved
1	WRITE	1=Write, 0=Read
16	REGISTER_ADD	Register Address
N*8	WRITE_DATA	Write Data; ignored for Reads

The DEVICE_ID is used to address PHY devices which are integrated into the same physical package and share a single SPI select. DEVICE_ID MUST be zero for accesses to single PHY devices.

Annex G (normative): Compatibility with Previous Versions of DOCSIS

DOCSIS 3.0 is the fourth generation of the DOCSIS specification. The terms DOCSIS 3.0, DOCSIS 2.0, DOCSIS 1.1 and DOCSIS 1.0 refer to these four different specifications.

The DOCSIS 3.0 specification primarily increases upstream and downstream throughput through the use of channel bonding, enhances security, adds enhanced support for multicast services and adds support for IPv6.

As well as supporting DOCSIS 3.0 CMs, the DOCSIS 3.0 CMTS MUST interoperate seamlessly with DOCSIS 2.0, DOCSIS 1.1 and DOCSIS 1.0 CMs. Furthermore, DOCSIS 3.0 CMs MUST interoperate seamlessly with DOCSIS 2.0, DOCSIS 1.1 and DOCSIS 1.0 CMTSs. Therefore, it is necessary for a DOCSIS 3.0 CM to function like a 1.0 CM when interoperating with a 1.0 CMTS, to function like a 1.1 CM when interoperating with a 1.1 CMTS and to function like a 2.0 CM when interoperating with a 2.0 CMTS.

This clause describes the interoperability issues and trade-offs involved when the operator wishes to support DOCSIS 2.0, DOCSIS 1.1 and/or DOCSIS 1.0 CMs as well as DOCSIS 3.0 CMs on the same cable access channel.

G.1 General Interoperability Issues

This clause addresses the general DOCSIS 1.x/2.0/3.0 interoperability issues that do not depend on the modulation type used for the upstream channel.

G.1.1 Initial Ranging

A DOCSIS CM's first upstream transmission is a ranging request message. This message may be a B-INIT-RNG-REQ, an INIT-RNG-REQ or a RNG-REQ depending on the CM's version, the type of channel on which the CM is ranging and the presence of the MDD. Refer to table 6-28 lists the type of message used under the different situations by modems capable of supporting downstream channel bonding.

DOCSIS 2.0 CMs performing initial ranging on a type 3 upstream transmit the INIT-RNG-REQ while 2.0 CMs ranging on a type 1 or 2 upstream and all DOCSIS 1.x CMs transmit the RNG-REQ.

G.1.2 Topology Resolution

DOCSIS 3.0 supports upstream and downstream topology resolution. DOCSIS 3.0 CMTSs SHOULD attempt topology resolution on pre-3.0 DOCSIS CMs. To aid in downstream topology resolution, DOCSIS 3.0 adds a downstream channel list to the MDD message. CMs supporting this message attempt to acquire downstream channels from the list and report back the resolution in the B-INIT-RNG-REQ. To aid in upstream topology resolution, DOCSIS 3.0 adds an Upstream Channel Adjustment TLV to the RNG-RSP that allows the CMTS to instruct a CM to move to a different upstream channel without the re-initialization that would be required with an upstream channel override in the RNG-RSP. This Upstream Channel Adjustment TLV is only applicable when a CM has transmitted a B-INIT-RNG-REQ.

For those modems not transmitting a B-INIT-RNG-REQ, the downstream frequency override in the RNG-RSP can be used to force the CM to attempt acquisition of a new downstream channel. Similarly, the upstream channel override portion of the RNG-RSP can be used to force the CM to attempt ranging on a new upstream channel prior to registration. The use of the upstream channel override in the RNG-RSP will result in the CM beginning initial ranging on the new upstream channel. Refer to clause 6.4.6.4 or its 3.0 equivalent.

G.1.3 Early Authentication and Encryption (EAE)

DOCSIS 3.0 supports early authentication and encryption. A CMTS advertises this capability in the MDD message. When a DOCSIS 3.0 CM sees an MDD enabling early authentication and encryption, the CM attempts to perform EAE per the [15] after ranging and ambiguity resolution. If the CM does not see an MDD enabling early authentication, then the CM does not initiate this process and moves on to establishing IP connectivity. Pre-3.0 DOCSIS CMs that do not support early authentication will not initiate this process. Modems not initiating EAE will initiate Baseline Privacy Initialization, if enabled in configuration file, after completing registration and prior to going operational.

G.1.4 Provisioning

The parameters of the TFTP configuration file for a DOCSIS 3.0 CM are a superset of those for Pre-3.0 DOCSIS CMs. The DOCSIS 3.0 configuration file contains 11 new top-level TLVs and many additional sub-fields to previously existing TLVs. The new top-level TLVs for configuration files are:

- SNMPv1v2c Coexistence.
- SNMPv3 Access View.
- SNMP CPE Access Control.
- Channel Assignment Configuration.
- CMTS Static Multicast Session.
- Software Upgrade IPv6 TFTP Server.
- TFTP Provisioned Modem IPv6 Address.
- Upstream Drop Classifier.
- Subscriber Mgmt CPE IPv6 Prefix List.
- Upstream Drop Classifier Group ID.
- Subscriber Mgmt Control Max CPE IPv6 Prefix.

Configuration file editors that support earlier versions of the DOCSIS specification may need to be modified to support these new TLVs and the new sub-fields added to support channel bonding and other features of DOCSIS 3.0.

A TFTP configuration file containing Class of Service TLVs is considered a "DOCSIS 1.0 style" configuration file. A TFTP configuration file containing Service Flow TLVs is considered a "DOCSIS 1.1/2.0/3.0 style" configuration file. A TFTP configuration file containing both Class of Service and Service Flow TLVs will be rejected by the CMTS (see clause 10.2.6.2).

If a DOCSIS 3.0 CM is provisioned with a DOCSIS 1.0-style TFTP configuration file, it will register as specified in the next clause, although in the REG-REQ or REG-REQ-MP it MUST still specify "DOCSIS 3.0" in the DOCSIS Version Modem Capability and MAY specify additional advanced (i.e. DOCSIS 1.1, DOCSIS 2.0 and DOCSIS 3.0) Modem Capabilities that it supports. Thus, a DOCSIS 3.0 CM can be provisioned to work seamlessly on a DOCSIS 1.0, a DOCSIS 1.1, a DOCSIS 2.0 or a DOCSIS 3.0 CMTS. However, a DOCSIS 3.0 modem on a pre-3.0 DOCSIS CMTS would be unable to support any DOCSIS 3.0-specific features not supported by that CMTS.

NOTE: DOCSIS 1.0 Class of Service encodings are incompatible with Multiple Transmit Channel mode operation, since configuration of Multiple Transmit Channel mode parameters (e.g. SID Clusters, Ranging SID) is tied to the existence of Service Flows.

Hence, the CMTS MUST disable Multiple Transmit Channel mode for a modem which includes DOCSIS 1.0 Class of Service encodings in its Registration Request. In contrast, Multiple Receive Channel mode does not depend on Service Flows, so the CMTS is permitted to enable Multiple Receive Channel mode on a modem using DOCSIS 1.0 Class of Service encodings.

A DOCSIS 3.0 CM operating on an S-CDMA channel with the Maximum Scheduled Codes feature enabled (see clause 10.2.6.2) and provisioned with a DOCSIS 1.0-style configuration file, SHOULD support fragmentation and indicate that support in the Modem Capabilities Encoding in the REG-REQ or REG-REQ-MP message. If a DOCSIS 3.0 CM supports certain advanced capabilities when registered as a DOCSIS 1.0 CM (as indicated by the Modem Capabilities Encoding), those features MUST function according to the requirements defined in the DOCSIS 3.0 specifications.

Consider the example of a DOCSIS 1.0 CM which does not recognize (and ignores) many of the new TLVs in a DOCSIS 1.1/2.0/3.0 style configuration file. This CM will be unable to register successfully if provisioned with a DOCSIS 1.1/2.0/3.0 configuration file. To prevent any functionality mismatches, a DOCSIS 3.0 CMTS MUST reject any Registration Request with DOCSIS 1.1/2.0/3.0-specific configuration parameters that are not supported by the associated Modem Capabilities encoding in the REG-REQ or REG-REQ-MP (see clause C.1.3.1).

A summary of TLV encodings is shown in table G-1.

Table G-1: Summary of TLV Encodings

Type	Description	First DOCSIS Version	Usage
0	Pad	1.0	Cfg File
1	Downstream Frequency	1.0	Cfg File, REG
2	Upstream Channel ID	1.0	Cfg File, REG
3	Network Access Control Object	1.0	Cfg File, REG
4	DOCSIS 1.0 Class of Service	1.0	Cfg File, REG
4.1	Class ID	1.0	Cfg File, REG
4.2	Maximum Downstream Rate	1.0	Cfg File, REG
4.3	Maximum Upstream Rate	1.0	Cfg File, REG
4.4	Upstream Channel Priority	1.0	Cfg File, REG
4.5	Guaranteed Minimum Upstream Channel Data Rate	1.0	Cfg File, REG
4.6	Maximum Upstream Channel Transmit Burst	1.0	Cfg File, REG
4.7	Class of Service Privacy Enable	1.0	Cfg File, REG
6	CM Message Integrity Check (MIC)	1.0	Cfg File, REG
7	CMTS Message Integrity Check (MIC)	1.0	Cfg File, REG
9	SW Upgrade Filename	1.0	Cfg File
10	SNMP Write Access Control	1.0	Cfg File
11	SNMP MIB Object	1.0	Cfg File
14	CPE Ethernet MAC Address	1.0	Cfg File
15	Telephone Settings Option (deprecated)	1.0	Cfg File
17	Baseline Privacy	1.0	Cfg File, REG
18	Max Number of CPEs	1.0	Cfg File, REG
19	TFTP Server Timestamp	1.0	Cfg File, REG
20	TFTP Server Provisioned Modem IPv4 Address	1.0	Cfg File, REG
21	SW Upgrade IPv4 TFTP Server	1.0	Cfg File
43	DOCSIS Extension Field/ (Vendor-Specific Vendor Encoding in 1.0)	1.0	Cfg File, REG
43.8	Reserved for Vendor ID Encoding (TLV 8)	1.0	Cfg File, REG
255	End-of-Data	1.0	Cfg File
22	Upstream Packet Classification	1.1	Cfg File, REG, DSx
23	Downstream Packet Classification	1.1	Cfg File, REG, DSx
24/25	Service Flow	1.1	Cfg File, REG, DSx
24/25.1	Service Flow Reference	1.1	Cfg File, REG
24/25.2	Service Flow Identifier	1.1	REG, DSx
24/25.3	Service Identifier	1.1	REG, DSx
24/25.4	Service Class Name	1.1	Cfg, REG, DSx
24/25.5	Service Flow Error Encodings	1.1	REG, DSx
24/25.5.1	Errored Parameter	1.1	REG, DSx
24/25.5.2	Error Code	1.1	REG, DSx
24/25.5.3	Error Message	1.1	REG, DSx
24/25.6	Quality of Service Parameter Set Type	1.1	Cfg File, REG, DSx
24/25.7	Traffic Priority	1.1	Cfg File, REG, DSx
24/25.9	Maximum Traffic Burst	1.1	Cfg File, REG, DSx
24/25.10	Minimum Reserved Traffic Rate	1.1	Cfg File, REG, DSx
24/25.11	Assumed Minimum Reserved Rate Packet Size	1.1	Cfg File, REG, DSx
24/25.12	Timeout for Active QoS Parameters	1.1	Cfg File, REG, DSx
24/25.13	Timeout for Admitted QoS Parameters	1.1	Cfg File, REG, DSx

Type	Description	First DOCSIS Version	Usage
24	Upstream Service Flow	1.1	Cfg File, REG, DSx
24.8	Upstream Maximum Sustained Traffic Rate	1.1	Cfg File, REG, DSx
24	Upstream Service Flow	1.1	Cfg File, REG, DSx
24.14	Maximum Concatenated Burst	1.1	Cfg File, REG, DSx
24.15	Service Flow Scheduling Type	1.1	Cfg File, REG, DSx
24.16	Request/Transmission Policy	1.1	Cfg File, REG, DSx
24.17	Nominal Polling Interval	1.1	Cfg File, REG, DSx
24.18	Tolerated Poll Jitter	1.1	Cfg File, REG, DSx
24.19	Unsolicited Grant Size	1.1	Cfg File, REG, DSx
24.20	Nominal Grant Interval	1.1	Cfg File, REG, DSx
24.21	Tolerated Grant Jitter	1.1	Cfg File, REG, DSx
24.22	Grants per Interval	1.1	Cfg File, REG, DSx
24.23	IP Type Of Service (DSCP) Overwrite	1.1	Cfg File, REG, DSx
24.24	Unsolicited Grant Time Reference	1.1	Cfg File, REG, DSx
25	Downstream Service Flow	1.1	Cfg File, REG, DSx
25.8	Downstream Maximum Sustained Traffic Rate	1.1	Cfg File, REG, DSx
25.14	Maximum Downstream Latency	1.1	Cfg File, REG, DSx
26	Payload Header Suppression	1.1	Cfg File, REG, DSx
26.1	Classifier Reference	1.1	Cfg File, REG, DSx
26.2	Classifier Identifier	1.1	REG, DSx
26.3	Service Flow Reference	1.1	Cfg File, REG, DSx
26.4	Service Flow Identifier	1.1	REG, DSx
26.43	Vendor Specific PHS Parameters	1.1	Cfg File, REG, DSx
26.7	Payload Header Suppression Field (PHSF)	1.1	Cfg File, REG, DSx
26.8	Payload Header Suppression Index (PHSI)	1.1	REG, DSx
26.9	Payload Header Suppression Mask (PHSM)	1.1	Cfg File, REG, DSx
26.10	Payload Header Suppression Size (PSSS)	1.1	Cfg File, REG, DSx, DBC
26.11	Payload Header Suppression Verification (PHSV)	1.1	Cfg File, REG, DSx, DBC
26.12	Reserved	-	-
28	Maximum Number of Classifiers	1.1	Cfg File, REG
29	Privacy Enable	1.1	Cfg File, REG
32	Manufacturer Code Verification Certificate	1.1	Cfg File
33	Co-Signer Code Verification Certificate	1.1	Cfg File
34	SNMPv3 Kickstart Value	1.1	Cfg File
34.1	SNMPv3 Kickstart Security Name	1.1	Cfg File
34.2	SNMPv3 Kickstart Mgr Public Num.	1.1	Cfg File
35	Subscriber Mgmt Control	1.1	Cfg File, REG
36	Subscriber Mgmt CPE IPv4 List	1.1	Cfg File, REG
37	Subscriber Mgmt Filter Groups	1.1	Cfg File, REG
38	SNMPv3 Notification Receiver	1.1	Cfg File
38.1	SNMPv3 Notification Rx IP Addr	1.1	Cfg File
38.2	SNMPv3 Notification Rx UDP port	1.1	Cfg File, REG
38.3	SNMPv3 Notification Rx Trap Type	1.1	Cfg File
38.4	SNMPv3 Notification Rx Timeout	1.1	Cfg File
38.5	SNMPv3 Notification Rx Retries	1.1	Cfg File
38.6	SNMPv3 Notification Rx Filtering Params	1.1	Cfg File
38.7	SNMPv3 Notification Rx Security Name	1.1	Cfg File
22/23/60.1	Classifier Reference	1.1	Cfg File, REG, DSx
22/23/60.2	Classifier Identifier	1.1	REG, DSx
22/23.3	Service Flow Reference	1.1	Cfg File, REG, DSx
22/23.4	Service Flow Identifier	1.1	REG, DSx
22/23/60.5	Rule Priority	1.1	Cfg File, REG, DSx
22/23	Classifier Activation State	1.1	Cfg File, REG, DSx
22/23/60.7	Dynamic Service Change Action	1.1	DSx
22/23/60.8	Classifier Error Encodings	1.1	REG, DSx
22/23/60.8.1	Errored Parameter	1.1	REG, DSx
22/23/60.8.2	Error Code	1.1	REG, DSx
22/23/60.8.3	Error Message	1.1	REG, DSx
22/23/60.9	IPv4 Packet Classification Encodings	1.1	Cfg File, REG, DSx
22/23/60.9.1	IPv4 Type of Service Range and Mask	1.1	Cfg File, REG, DSx
22/23/60.9.2	IP Protocol	1.1	Cfg File, REG, DSx
22/23/60.9.3	IPv4 Source Address	1.1	Cfg File, REG, DSx
22/23/60.9.4	IPv4 Source Mask	1.1	Cfg File, REG, DSx

Type	Description	First DOCSIS Version	Usage
22/23/60.9.5	IPv4 Destination Address	1.1	Cfg File, REG, DSx
22/23/60.9.6	IPv4 Destination Mask	1.1	Cfg File, REG, DSx
22/23/60.9.7	TCP/UDP Source Port Start	1.1	Cfg File, REG, DSx
22/23/60.9.8	TCP/UDP Source Port End	1.1	Cfg File, REG, DSx
22/23/60.9.9	TCP/UDP Destination Port Start	1.1	Cfg File, REG, DSx
22/23/60.9.10	TCP/UDP Destination Port End	1.1	Cfg File, REG, DSx
22/23/60.10	Ethernet LLC Packet Classification Encodings	1.1	Cfg File, REG, DSx
22/23/60.10.1	Destination MAC Address	1.1	Cfg File, REG, DSx
22/23/60.10.2	Source MAC Address	1.1	Cfg File, REG, DSx
22/23/60.10.3	Ethertype/DSAP/Mac Type	1.1	Cfg File, REG, DSx
22/23/60.11	IEEE 802.1P/Q Packet Classification Encodings	1.1	Cfg File, REG, DSx
22/23/60.11.1	IEEE 802.1P User Priority	1.1	Cfg File, REG, DSx
22/23/60.11.2	IEEE 802.1Q VLAN ID	1.1	Cfg File, REG, DSx
22/23/60.43	Vendor Specific Classifier Parameters	1.1	Cfg File, REG, DSx
26.6.1	Errored Parameter	1.1	REG, DSx
26.6.2	Error Code	1.1	REG, DSx
26.6.3	Error Message	1.1	REG, DSx
24/25.43	Vendor Specific QoS Parameters	2.0	Cfg File, REG, DSx
39	Enable 2.0 Mode	2.0	Cfg File, REG
40	Enable Test Modes	2.0	Cfg File, REG
41	Downstream Channel List	2.0	Cfg File, REG
41.1	Single DS Channel	2.0	Cfg File, REG
41.1.1	Single DS Chan Timeout	2.0	Cfg File, REG
41.1.2	Single DS Chan Frequency	2.0	Cfg File, REG
41.2	DS Frequency Range	2.0	Cfg File, REG
41.2.1	DS Freq. Range Timeout	2.0	Cfg File, REG
41.2.2	DS Frequency Range Start	2.0	Cfg File, REG
41.2.3	DS Frequency Range End	2.0	Cfg File, REG
41.2.4	DS Frequency Range Step Size	2.0	Cfg File, REG
41.3	Default Scanning	2.0	Cfg File, REG
42	Static Multicast MAC Address	2.0	Cfg File
43.1	CM Load Balancing Policy ID	2.0	Cfg File, REG
43.2	CM Load Balancing Priority	2.0	Cfg File, REG
43.3	CM Load Balancing Group ID	2.0	Cfg File, REG
43.4	CM Ranging Class ID Extension	2.0	Cfg File, REG
43.5	L2VPN Encoding	2.0	Cfg File, REG
45	Downstream Unencrypted Traffic (DUT) Filtering	2.0	Cfg File, REG
65	L2VPN MAC Aging Encoding	2.0	Cfg File
22/23/60.12	IPv6 Packet Classification Encodings	3.0	Cfg File, REG, DSx
22/23/60.12.1	IPv6 Traffic Class	3.0	Cfg File, REG, DSx
22/23/60.12.2	IPv6 Flow Label	3.0	Cfg File, REG, DSx
22/23/60.12.3	IPv6 Next Header Type	3.0	Cfg File, REG, DSx
22/23/60.12.4	IPv6 Source Address	3.0	Cfg File, REG, DSx
22/23/60.12.5	IPv6 Source Prefix Length (bits)	3.0	Cfg File, REG, DSx
22/23/60.12.6	IPv6 Destination Address	3.0	Cfg File, REG, DSx
22/23/60.12.7	IPv6 Destination Prefix Length (bits)	3.0	Cfg File, REG, DSx
22/60.13	CM Interface Mask (CMIM)	3.0	Cfg File, REG, DSx
24/25.31	Service Flow Required Attribute Mask	3	Cfg File, REG, DSx
24/25.32	Service Flow Forbidden Attribute Mask	3	Cfg File, REG, DSx
24/25.33	Service Flow Attribute Aggregation Rule Mask	3	Cfg File, REG, DSx
24/25.34	Application Identifier	3	Cfg File, REG, DSx
24.25	Multiplier to Contention Request Backoff Window	3.0	REG, DSx
24.26	Multiplier to Number of Bytes Requested	3.0	Cfg File, REG, DSx
24/25.27	Peak Traffic Rate	3.0	Cfg File, REG, DSx
25.15	Reserved	-	-
25.23	IP Type Of Service (DSCP) Overwrite	3.0	Cfg File, REG, DSx
25.17	Downstream Resequencing	3	Cfg File, REG, DBC
38.8	SNMPv3 Notification Receiver IPv6 Address	3	Cfg File
43.6	Extended CMTS MIC config	3.0	Cfg File, REG
43.6.1	Extended CMTS MIC HMAC type	3.0	Cfg File, REG
43.6.2	Extended CMTS MIC Bitmap	3	Cfg File, REG
43.6.3	Explicit Extended CMTS MIC Digest Subtype	3.0	Cfg File, REG
43.7	SAV Authorization Encoding	3.0	Cfg File, REG

Type	Description	First DOCSIS Version	Usage
43.7.1	SAV Group Name	3.0	Cfg File, REG
43.7.2	SAV Static Prefix	3.0	Cfg File, REG
43.9	CM Attribute Masks	3.0	Cfg File, REG
43.9.1	CM Required Downstream Attribute Mask	3.0	Cfg File, REG
43.9.2	CM Downstream Forbidden Attribute Mask	3.0	Cfg File, REG
43.9.3	CM Upstream Required Attribute Mask	3.0	Cfg File, REG
43.9.4	CM Upstream Forbidden Attribute Mask	3.0	Cfg File, REG
43.10	IP Multicast Join Authorization	3.0	Cfg File, REG
43.10.1	IP Multicast Profile Name	3.0	Cfg File, REG
43.10.2	IP Multicast Join Authorization Static Session Rule	3.0	Cfg File, REG
43.10.3	Maximum Multicast Sessions	3.0	Cfg File, REG
43.11	Service Type identifier	3	Cfg File, REG
53	SNMPv1v2c Coexistence	3.0	Cfg File
53.1	SNMPv1v2c Community Name	3.0	Cfg File
53.2	SNMPv1v2c Transport Address Access	3.0	Cfg File
53.2.1	SNMPv1v2c Transport Address	3.0	Cfg File
53.2.2	SNMPv1v2c Transport Address Mask	3.0	Cfg File
53.3	SNMPv1v2c Access View Type	3.0	Cfg File
53.4	SNMPv1v2c Access View Name	3.0	Cfg File
54	SNMPv3 Access View	3.0	Cfg File
54.1	SNMPv3 Access View Name	3.0	Cfg File
54.2	SNMPv3 Access View Subtree	3.0	Cfg File
54.3	SNMPv3 Access View Mask	3.0	Cfg File
54.4	SNMPv3 Access View Type	3.0	Cfg File
55	SNMP CPE Access Control	3.0	Cfg File
56	Channel Assignment Configuration Settings	3.0	Cfg File, REG
56.1	Transmit Channel Assignment	3.0	Cfg File, REG
56.2	Receive Channel Assignment	3.0	Cfg File, REG
58	Software Upgrade IPv6 TFTP Server	3.0	Cfg File
59	TFTP Server Provisioned Modem IPv6 Address	3.0	Cfg File, REG
60	Upstream Drop Packet Classification	3.0	Cfg File, REG, DSC
61	Subscriber Mgmt CPE IPv6 Prefix List	3.0	Cfg File, REG
62	Upstream Drop Classifier Group ID	3.0	Cfg File, REG
63	Subscriber Mgmt Control Max CPE IPv6 Prefix	3.0	Cfg File, REG
64	CMTS Static Multicast Session Encoding	3.0	Cfg File
64.1	Static Multicast Group Encoding	3.0	Cfg File
64.2	Static Multicast Source Encoding	3.0	Cfg File
64.3	Static Multicast CMIM Encoding	3.0	Cfg File
201,202, 216-231	eSAFE Configuration	[7]	Cfg File
66	Management Event Control Encoding	3.0	Cfg File

G.1.5 Registration

The registration procedure specified in DOCSIS 3.0 is very different from earlier versions of the specification. DOCSIS 3.0-style registration provides for resolution of the CM's upstream and downstream service groups and the provisioning of multiple downstream and upstream channels. The CMTS announces its support for the 3.0-style registration by transmitting an MDD on the downstream channel. When a CM supporting DOCSIS 3.0 style registration initializes, it acquires a downstream and looks for SYNC messages. If the CM finds SYNC messages and an MDD message on the downstream, it attempts to resolve downstream ambiguity using any hints supplied by the MDD.

When the CM sends a REG-REQ or REG-REQ-MP message, it includes TLVs relating the new capabilities added as part of DOCSIS 3.0. Should the CM find SYNC messages on a downstream channel that does not contain an MDD, then the CM uses DOCSIS2.0-style registration but includes its 3.0 modem capabilities.

For CMTSs and CMs that do not support the DOCSIS 3.0-style registration, registration will occur as described below.

A DOCSIS 3.0 CMTS is designed to handle the registration TLVs from DOCSIS 1.0 CMs as well as the TLVs introduced in DOCSIS 1.1 (TLV types 22 to 38), DOCSIS 2.0 (TLV types 39 to 42 and 45) and DOCSIS 3.0 (TLV types 46 to 52, 56 and 57). Furthermore a DOCSIS 3.0 CM can handle any TLVs in a configuration file usable by a DOCSIS 1.0 CM.

A DOCSIS 1.1, 2.0 or 3.0 CM could be configured to use the Service Class Name which is statically defined at the CMTS instead of explicitly asking for the service class parameters. When the DOCSIS 3.0 CMTS receives such a Registration-Request, it encodes the actual parameters of that service class in the Registration-Response and expects the 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 (or any CM using a 1.0-style configuration file) registers with the same CMTS, the absence of Service Class Names eliminates the need for the DOCSIS 3.0 CMTS to explicitly specify the service class parameters in the Registration-Response using DOCSIS 1.1, 2.0 or 3.0 TLVs. The Registration-Request from a DOCSIS 1.0 CM explicitly requests all non-default service class parameters in the Registration-Request per the CM's provisioning information. When a DOCSIS 3.0 CMTS receives a Registration-Request containing DOCSIS 1.0 Class of Service Encodings, it will respond with the DOCSIS 1.0-style Registration-Response and, if the CM is a DOCSIS 1.x CM or is operating on a Type 1 channel, not expect the CM to send the Registration-Acknowledge MAC message. A DOCSIS 1.0 CM can be further identified by the absence of the "DOCSIS Version" Modem Capabilities encoding in the Registration-Request.

In the case where a DOCSIS 2.0 CM or a DOCSIS 3.0 CM using DOCSIS 2.0-style registration is using a DOCSIS 1.0-style configuration file, there is an additional consideration. This is because, in the case where the upstream is a Type 2 or Type 4 upstream (see clause 6.4.3) and therefore supports both TDMA and A-TDMA features, the Registration-Acknowledge message is also used to synchronize switching from TDMA (DOCSIS 1.x) operation to A-TDMA (DOCSIS 2.0) operation. It is important that this switch be coordinated correctly between the CM and the CMTS in order for the CMTS to be able to correctly interpret bandwidth requests from the CM (see clause 10.2.6). Therefore, when a DOCSIS 2.0 or 3.0 CM registers using a 1.0-style configuration file with Enable 2.0 Mode enabled on a Type 2, Type 3 or Type 4 upstream, it transmits a Registration-Acknowledgment with a confirmation code of OK/SUCCESS (since 1.0-style registration does not allow for the CM to reject the Registration-Response). The CMTS knows to expect this because the modem capabilities field in the Registration-Request indicated that the CM was a 2.0 or 3.0 CM.

A DOCSIS 3.0 CM using DOCSIS 3.0-style registration will always send a REG-ACK upon receiving a REG-RSP or REG-RSP-MP, regardless of the DOCSIS version of the configuration file.

Table G-2 summarizes registration behavior for all cases involving a DOCSIS 3.0 CM.

Table G-2: Registration Acknowledgement Behavior for a DOCSIS 3.0 CM

Configuration file	Type 1 Upstream, no MDDs on Downstream	Type 2 Upstream, no MDDs on Downstream	Type 3 or 4 Upstream, no MDDs on Downstream	Type 1 or 2 Upstream, MDDs present on Downstream	Type 3 or 4 Upstream, MDDs present on Downstream
1.0-style configuration file that disables DOCSIS 2.0 mode	CM does not send REG-ACK.	CM does not send REG-ACK.	N/A - CM does not attempt registration clause 10.2.5.7	CM sends REG-ACK.	N/A - CM does not attempt registration clause 10.2.5.7
1.0-style configuration file that does not disable DOCSIS 2.0 mode	CM does not send REG-ACK.	CM sends REG-ACK with SUCCESS confirmation code. CMTS waits for REG-ACK.	CM sends REG-ACK with SUCCESS confirmation code. CMTS waits for REG-ACK.	CM sends REG-ACK.	CM sends REG-ACK.
1.1/2.0/3.0-style configuration file	CM sends REG-ACK.	CM sends REG-ACK.	CM sends REG-ACK.	CM sends REG-ACK.	CM sends REG-ACK.

Table G-3 shows the registration operation of the various versions of DOCSIS CMs with a 1.0-style configuration file registering on the various upstream channel types.

Table G-3: Registration Operation of DOCSIS CMs with 1.0-style Config File

	Type 1 Channel	Type 2 Channel	Type 3 Channel	Type 4 Channel
1.0 CM	No REG-ACK	No REG-ACK	N/A	N/A
1.1 CM	No REG-ACK	No REG-ACK	N/A	N/A
2.0 CM with DOCSIS 2.0 operation disabled in config file	No REG-ACK	No REG-ACK	No REG-ACK	N/A
2.0 CM with DOCSIS 2.0 operation not disabled in config file	No REG-ACK	REG-ACK	REG-ACK	N/A
3.0 CM with DOCSIS 2.0 operation disabled in config file, no MDD message	No REG-ACK	No REG-ACK	No REG-ACK	No REG-ACK
3.0 CM with DOCSIS 2.0 operation not disabled in config file, no MDD message	No REG-ACK	REG-ACK	REG-ACK	REG-ACK
3.0 CM, MDD message present	REG-ACK	REG-ACK	REG-ACK	REG-ACK

Another minor issue 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 a DOCSIS 3.0 CMTS typically operates with unidirectional service classes, it can easily translate a DOCSIS 1.0 Class-of-Service Configuration Setting into DOCSIS 1.1, 2.0 or 3.0 Service Flow Encodings for setting up unidirectional service classes in local QoS implementation. However, for DOCSIS 1.0 modems, the DOCSIS 3.0 CMTS continues to maintain the QoSProfile table (with bi-directional Class parameters) for backward compatibility with the DOCSIS 1.0 MIB.

Thus, if properly provisioned, a DOCSIS 1.0, a DOCSIS 1.1, a DOCSIS 2.0 and a DOCSIS 3.0 CM can all successfully register with the same DOCSIS 3.0 CMTS and a DOCSIS 3.0 CM can register with a 1.0 CMTS. Furthermore, a DOCSIS 3.0 CM can use a DOCSIS 1.0-style configuration file, register on a DOCSIS 3.0 CMTS and still use DOCSIS 2.0 and DOCSIS 3.0 enhanced physical-layer features with DOCSIS 1.0 class-of-service features.

Table G-4 shows the registration parameters that cannot be included in the configuration file.

Table G-4: Summary of Registration Parameters not in Configuration File

Type	Description	First DOCSIS Version	Usage
5	Modem Capabilities	1.0	REG
5.1	Concatenation Support	1.0	REG
8	Vendor ID Encoding	1.0	REG
12	Modem IP Address	1.0	REG
13	Service(s) Not Available Response	1.0	REG
5.2	DOCSIS Version	1.1	REG
5.3	Fragmentation Support	1.1	REG
5.4	PHS Support	1.1	REG
5.5	IGMP Support	1.1	REG
5.6	Privacy Support	1.1	REG
5.7	Downstream SAID Support	1.1	REG
5.8	Upstream Service Flow Support	1.1	REG
5.9	Optional Filtering Support	1.1	REG
5.10	Transmit Equalizer Taps per Modulation Int.	1.1	REG
5.11	Number of Transmit Equalizer Taps	1.1	REG
5.12	DCC Support	1.1	REG
27	HMAC-Digest	1.1	DSx, DBC
30	Authorization Block	1.1	DSx
31	Key Sequence Number	1.1	DSx, DBC
5.13	IP Filters Support	2.0	REG
5.14	LLC Filters Support	2.0	REG
5.15	Expanded Unicast SID Space	2.0	REG
5.16	Ranging Hold-Off Support	2.0	REG
5.17	L2VPN Capability	2.0	REG
5.18	L2VPN eSAFE Host Capability	2.0	REG
5.19	DS Unencrypted Traffic (DUT) Filtering	2.0	REG
5.20	Upstream Frequency Range Support	3.0	REG
5.21	Upstream Symbol Rate Support	3.0	REG
5.22	SAC Mode 2 Support	3.0	REG
5.23	Code Hopping Mode 2 Support	3.0	REG
5.24	Multiple Transmit Channel Support	3.0	REG
5.25	5.12 Msps US Transmit Channel Support	3.0	REG
5.26	2.56 Msps US Transmit Channel Support	3.0	REG
5.27	Total SID Cluster Support	3.0	REG
5.28	SID Clusters per Service Flow Support	3.0	REG
5.29	Multiple Receive Channel Support	3.0	REG
5.30	Total DS Service ID (DSID) Support	3.0	REG
5.31	Resequencing DSID Support	3.0	REG
5.32	Multicast Downstream SID (DSID) Support	3.0	REG
5.33	Multicast DSID Forwarding	3.0	REG
5.34	Frame Control Type Forwarding Capability	3.0	REG
5.35	DPV Capability	3.0	REG
5.36	Unsolicited Grant Service US SF Support	3.0	REG
5.37	MAP and UCD Receipt Support	3.0	REG
5.38	Upstream Drop Classifier Support	3.0	REG
5.39	IPv6 Support	3.0	REG
44	Vendor Specific Capabilities	2.0	REG
46	Transmit Channel Config	3.0	REG, DBC
46.1	TCC Reference	3.0	REG, DBC
46.2	Upstream Channel Action	3.0	REG, DBC
46.3	Upstream Channel ID	3.0	REG, DBC
46.4	New Upstream Channel ID	3.0	REG, DBC
46.5	UCD	3.0	REG, DBC
46.6	Ranging SID	3.0	REG, DBC
46.7	Initialization Technique	3.0	REG, DBC
46.8	Ranging Parameters	3.0	REG, DBC
46.8.1	Ranging Reference Channel ID	3.0	REG, DBC
46.8.2	Timing Offset, Integer Part	3.0	REG, DBC
46.8.3	Timing Offset, Fractional Part	3.0	REG, DBC
46.8.4	Power Offset	3.0	REG, DBC
46.254	TCC Error Encodings	3.0	REG, DBC
46.254.1	Reported Parameter	3.0	REG, DBC

Type	Description	First DOCSIS Version	Usage
46.254.2	Error Code	3.0	REG, DBC
46.254.3	Error Message	3.0	REG, DBC
47	Service Flow SID Cluster Assignment	3.0	REG, DSx, DBC
47.1	SFID	3.0	REG, DSx, DBC
47.2	SID Cluster Encoding	3.0	REG, DSx, DBC
47.2.1	SID Cluster ID	3.0	REG, DSx, DBC
47.2.2	SID-to-Channel Mapping	3.0	REG, DSx, DBC
47.3	SID Cluster Switchover Criteria	3.0	REG, DSx, DBC
47.3.1	Maximum Requests per SID Cluster	3.0	REG, DSx, DBC
47.3.2	Maximum Outstanding Bytes per SID Cluster	3.0	REG, DSx, DBC
47.3.3	Maximum Total Bytes Requested per SID Cluster	3.0	REG, DSx, DBC
47.3.4	Maximum Time in the SID Cluster	3.0	REG, DSx, DBC
48	Receive Channel Profile	3.0	REG
48.1	RCP ID (OUI + Profile)	3.0	REG
48.2	RCP Name	3.0	REG
48.3	RCP Center Frequency Spacing	3.0	REG
48.4	Receive Module Capability	3.0	REG
48.4.1	Receive Module Index (being described)	3.0	REG
48.4.2	Receive Module Adjacent Channels	3.0	REG
48.4.3	Receive Module Channel Block Range	3.0	REG
48.4.3.1	Receive Module Min Center Frequency	3.0	REG
48.4.3.2	Receive Module Max Center Frequency	3.0	REG
48.4.5	Receive Module Resequencing Chan. Sub.	3.0	REG
48.4.6	Receive Module Connectivity (descr.)	3.0	REG
48.4.7	Receive Module Common PHY Params	3.0	REG
48.5	Receive Channels (capability)	3.0	REG
48.5.1	Receive Channel Index (within RCP)	3.0	REG
48.5.2	Receive Channel Connectivity (Capability)	3.0	REG
48.5.3	Receive Channel Connected Offset	3.0	REG
48.5.5	Receive Channel Primary DS Chan Indic	3.0	REG
48.43	Receive Channel Profile Vendor Specific Parameters	3.0	REG
49	Receive Channel Config	3.0	REG, DBC
49.1	RCP-ID	3.0	REG, DBC
49.4	Receive Module Assignment	3.0	REG, DBC
49.4.1	Receive Module Index (being assigned)	3.0	REG, DBC
49.4.4	Receive Module First Channel Center Freq.	3.0	REG, DBC
49.4.6	Receive Module Connectivity (assigned)	3.0	REG, DBC
49.5	Receive Channels (assigned)	3.0	REG, DBC
49.5.1	Receive Channel Index (within RCC)	3.0	REG, DBC
49.5.2	Receive Channel Connectivity (Assigned)	3.0	REG, DBC
49.5.4	Receive Channel Center Freq. Assignment	3.0	REG, DBC
49.5.5	Receive Channel Primary DS Chan Indic	3.0	REG, DBC
49.6	Partial Service Downstream Channels	3.0	REG, DBC
49.43	Receive Channel Configuration Vendor Specific Parameters	3.0	REG, DBC
49.254	RCC Error Encodings	3.0	REG, DBC
49.254.1	Receive Module or Receive Channel	3.0	REG, DBC
49.254.2	Receive Module Index or Receive Channel Index	3.0	REG, DBC
49.254.3	Reported Parameter	3.0	REG, DBC
49.254.4	Error Code	3.0	REG, DBC
49.254.5	Error Message	3.0	REG, DBC
50	DSID Encodings	3.0	REG, DBC
50.1	Downstream Service Identifier	3.0	REG, DBC
50.2	Downstream Service Identifier Action	3.0	REG, DBC
50.3	Downstream Resequencing Encodings	3.0	REG, DBC
50.3.1	Resequencing DSID	3.0	REG, DBC
50.3.2	Downstream Resequencing Channel List	3.0	REG, DBC
50.3.3	DSID Resequencing Wait Time	3.0	REG, DBC
50.3.4	Resequencing Warning Threshold	3.0	REG, DBC
50.3.5	CM-STATUS Hold-Off Timer (Out of Rng)	3.0	REG, DBC
50.4	Multicast Encodings	3.0	REG, DBC
50.4.1	Client MAC Address Encodings	3.0	REG, DBC
50.4.1.1	Client MAC Address Action	3.0	REG, DBC
50.4.1.2	Client MAC Address	3.0	REG, DBC
50.4.2	Multicast CM Interface Mask	3.0	REG, DBC

Type	Description	First DOCSIS Version	Usage
50.4.3	Multicast Group MAC Addresses Encodings	3.0	REG, DBC
50.4.26.x	Payload Header Suppression Encodings	3.0	REG, DBC
51	Security Association Encoding	3.0	REG, DBC
51.1	SA Action	3.0	REG, DBC
51.23	SA-Descriptor	3.0	REG, DBC
52	Initializing Channel Timeout	3.0	REG, DBC

G.1.6 Requesting Bandwidth

All versions of DOCSIS CMs use mini-slot based requests (via Request Frame, REQ_EHDR or BPI EHDR) to request bandwidth prior to receiving the REG-RSP or REG-RSP-MP. If a CM receives a REG-RSP or REG-RSP-MP enabling Multiple Transmit Channel Mode, the CM immediately begins using queue-depth based requesting for all subsequent bandwidth requests. If the CM receives a REG-RSP or REG-RSP-MP disabling Multiple Transmit Channel Mode or if the CM did not previously advertise its ability to support Multiple Transmit Channel Mode, then the CM continues using mini-slot based requesting. The CMTS knows what type of requesting the CM is using based on the request format itself and the mode of operation it relayed to the CM during registration.

G.1.7 Encryption Support

The CM and CMTS may perform a Baseline Privacy Message exchange (either as part of Early Authentication and Encryption or as part of Baseline Privacy Initialization after registration). This message exchange includes an encryption suite exchange to ensure that the CMTS becomes aware of the supported cryptographic suites. The CMTS will not enable an encryption suite that the CM does not support.

G.1.8 Downstream Channel Bonding

Through the Multiple Receive Channel Support capability encoding in the REG-REQ or REG-REQ-MP, a CM informs the CMTS of the modem's ability to support downstream channel bonding. A DOCSIS 3.0 CMTS MUST NOT send a REG-RSP or REG-RSP-MP with a Receive Channel Configuration to a CM that has not advertised support of Multiple Receive Channels in the modem capability portion of the REG-REQ or REG-REQ-MP. If the CM does not include the Multiple Receive Channel Support capability encoding in the REG-REQ or REG-REQ-MP, then the CM is incapable of supporting Multiple Receive Channels.

G.1.9 Upstream Channel Bonding and Transmit Channel Configuration Support

Through the Multiple Transmit Channel Support modem capability encoding in the REG-REQ or REG-REQ-MP, a CM informs the CMTS of the modem's ability to support Multiple Transmit Channel Mode and/or the Transmit Channel Configuration (TCC). If the CM reports a Multiple Transmit Channel Support capability of zero, the CM is incapable of supporting Multiple Transmit Channel Mode, but is capable of understanding the TCC for a single channel in the REG-RSP or REG-RSP-MP and in a DBC-REQ. The CMTS MAY send a TCC in the REG-RSP or REG-RSP-MP to such a CM. If the CM reports a Multiple Transmit Channel Mode of one or greater, the CM is capable of supporting Multiple Transmit Channel Mode. The CMTS MAY enable Multiple Transmit Channel Mode through the REG-RSP or REG-RSP-MP. Should the CMTS choose to enable Multiple Transmit Channel Mode, the CMTS MUST include a TCC in the REG-RSP or REG-RSP-MP and use DBC messaging for upstream channel changes, even if only a single channel is being configured. The CMTS MUST NOT send a Multiple Transmit Channel Mode enable setting to a CM that did not include a non-zero Multiple Transmit Channel Support capability in the REG-REQ or REG-REQ-MP. Similarly, the CMTS MUST NOT send a Transmit Channel Configuration encoding in the REG-RSP or REG-RSP-MP to a CM that did not include the Multiple Transmit Channel Support capability (regardless of the value of that capability) in the REG-REQ or REG-REQ-MP.

Whenever the CMTS sends a TCC to a CM, the CMTS MUST use either DCC messaging, with an initialization technique of zero (re-initialize MAC) or DBC messaging to make any upstream channel changes.

G.1.10 Dynamic Service Establishment

There are 8 MAC messages that relate to Dynamic Service Establishment. A DOCSIS 1.0 CM will never send dynamic service messages since they are not supported. A DOCSIS 1.1, 2.0 or 3.0 CM will never send these messages to a DOCSIS 1.0 CMTS because in order to register successfully, the CM has to be provisioned as a DOCSIS 1.0 CM and will act accordingly. When a DOCSIS 1.1, 2.0 or 3.0 CM is connected to a DOCSIS 1.1, 2.0 or 3.0 CMTS, these dynamic service messages work as expected.

G.1.11 Fragmentation

Fragmentation is initiated by the CMTS. There are two styles of fragmentation. The first is the fragmentation introduced in DOCSIS 1.1. This type of fragmentation is controlled by the fragmentation modem capability encoding. Thus, a DOCSIS 1.0 CMTS will never initiate fragmentation since it knows nothing about it. A DOCSIS 1.1, 2.0 or 3.0 CMTS can only initiate this type of fragmentation for DOCSIS 1.1, 2.0 or 3.0 CMs. A DOCSIS 3.0 CMTS **MUST NOT** attempt to fragment transmissions from a CM that has not indicated a Modem Capabilities encoding for Fragmentation Support with a value of 1.

The second style of fragmentation is the continuous concatenation and fragmentation that is part of Multiple Transmit Channel Mode's segmentation introduced in DOCSIS 3.0. This type of fragmentation is linked to the Multiple Transmit Channel Support capability. If the CM reports a value greater than zero for this capability, the CMTS may enable this mode of fragmentation by returning a non-zero value. The CM will not use the first style of fragmentation once Multiple Transmit Channel Mode is enabled. The CMTS will not enable Multiple Transmit Channel Mode (including continuous concatenation and fragmentation) for a CM that has not reported support for this capability.

G.1.12 Multicast Support

It is mandatory for DOCSIS 1.0 CMs to support forwarding of multicast traffic. However, the specification is silent on IGMP support. 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 CMs on a DOCSIS 3.0 network. Multicast forwarding in DOCSIS 3.0 CMs is controlled by the Multicast DSID Forwarding capability exchange in all cases. Additional information on backward compatibility for multicast forwarding may be found in clause G.4.

G.1.13 Changing Upstream Channels

There are three mechanisms for changing an upstream channel after registration: DBC messaging, DCC messaging and UCC messaging. The message type used for changing an upstream channel depends on the CM and CMTS.

DBC messaging was introduced in DOCSIS 3.0 and can be used to change multiple upstream channels and multiple downstream channels simultaneously within a single MAC domain. This messaging includes an initialization technique that allows the CMTS to instruct the CM to do a specific type of ranging (or none at all) before transmitting data on the new upstream channel. DBC also allows the CMTS to give relative ranging adjustments to the new channel based on the ranging parameters of another channel assigned to the CM. This relative adjustment allows the CM to use known channel similarities in the ranging adjustment. The CMTS **MUST** support the use of DBC messaging to change channels whenever Multiple Transmit Channel Mode is enabled at the CM. If Multiple Transmit Channel Mode is not enabled but a Transmit Channel Configuration was assigned during registration, the CMTS **MUST** support the use of DBC messaging to switch the upstream channel of the CM.

DCC messaging was introduced in the DOCSIS 1.1. DCC messaging supports changing a single upstream channel when a CM is not operating in Multiple Transmit Channel Mode and a Transmit Channel Configuration was not assigned during registration. DCC messaging also supports moving the CM to a new MAC domain (with an initialization technique of re-initialize MAC) when the CM is operating in Multiple Transmit Channel Mode. Like DBC, DCC messaging allows the CMTS to change both upstream and downstream channels simultaneously and allows the CMTS to specify an initialization technique for the new upstream. The DCC messaging does not support the relative adjustments included in the DBC messaging. DCC messaging **MUST NOT** be used for upstream channel changes (other than changes between MAC domains) when Multiple Transmit Channel Mode is enabled for the CM. If Multiple Transmit Channel Mode is not enabled and a Transmit Channel Configuration was not assigned during registration, the CMTS **MAY** use DCC to switch the upstream channel of the CM.

For DOCSIS 1.0 CMs, the CMTS can only use UCC messaging to change an upstream channel. UCC messaging was introduced in DOCSIS 1.0 and provides a simple, though loosely controlled, mechanism for changing a single upstream channel. The CM receiving the UCC ranges on the new channel first using broadcast ranging opportunities.

G.1.14 Changing Downstream Channels

There are two mechanisms for changing downstream channels at a CM after registration: DBC messaging and DCC messaging. Both mechanisms allow simultaneous changing of upstream and downstream channels, but the DBC messaging is designed for multi-channel support. For a CM operating in Multiple Receive Channel Mode, the CMTS uses DBC messaging for changing downstream channels at that CM unless the CM is moving to another MAC domain, in which case DCC messaging can be used. To change a downstream channel for a CM not operating in Multiple Receive Channel Mode, the CMTS MUST NOT use DBC messaging. For a DOCSIS 1.1, 2.0 or 3.0 CM not operating in Multiple Receive Channel Mode, the CMTS should use DCC messaging to effect downstream channel changes.

For DOCSIS 1.0 CMs, the only mechanism to move the CM to a new downstream channel is to force a re-initialization of the CM.

G.2 Support for 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 CMTSs 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 CMs 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 DOCSIS 1.1. Unless a CM is fully DOCSIS 2.0 compliant, it MUST NOT send a "DOCSIS Version" Modem Capability which indicates DOCSIS 2.0. Similarly, unless a CM is fully DOCSIS 3.0 compliant, it MUST NOT send a "DOCSIS Version" Modem Capability which indicates DOCSIS 3.0.

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 xxxxxxxx MUST be an ASCII representation of the hexadecimal encoding of the Modem Capabilities. Refer to clauses C.1.3.1 and D.1.1 for details. 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 3.0 CMTS receives a 1.0-style Registration Request message (with CoS configuration settings) from a CM, the CMTS MUST, by default, force the modem to operate in a DOCSIS 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-2. The CMTS MAY provide switches, as indicated in table G-2, for the operator to selectively allow certain hybrid features to be enabled. As an exception to these defaults, the DOCSIS 3.0 CMTS SHOULD allow the use of fragmentation for DOCSIS 2.0 or 3.0 CMs registering in DOCSIS 1.0 mode on an S-CDMA channel that has the Maximum Scheduled Codes feature (see clause 8.3.1) enabled.

Table G-5: Hybrid Mode Controls

	Concatenation Support	Fragmentation Support	Privacy Support
1.0 CM	allow/[deny]	allow/[deny]	allow BPI+/[force BPI]
1.1 ohnr 2.0 CM with CoS	allow/[deny]	allow/[deny]	allow BPI+/[force BPI]
3.0 CM with CoS	allow/[deny]	allow/[deny]	allow BPI+/[force BPI]*

NOTE: If Early Authentication Encryption (EAE) is enabled in the MDD, a 3.0 CM is not forced to be in BPI mode.

A DOCSIS 2.0 hybrid CMTS (i.e. that supports features beyond DOCSIS 2.0 that are defined in DOCSIS 3.0) MAY leave supported Modem Capabilities defined in DOCSIS 3.0 set to "On" in the Registration Response. However, unless a CMTS is fully DOCSIS 3.0-compliant (i.e. not a hybrid), it MUST still set all "DOCSIS Version" Modem Capabilities to DOCSIS 2.0.

As always, any Modem Capability set to "Off" in the Registration Response is viewed as unsupported by the CMTS and MUST NOT be used by the CM.

G.3 Upstream Physical Layer Interoperability

G.3.1 DOCSIS 2.0 TDMA Interoperability

G.3.1.1 Mixed-mode operation with TDMA on a Type 2 channel

In mixed-mode operation with both DOCSIS 1.x and DOCSIS 2.0 TDMA, a single channel is defined with a single UCD that contains both type 4 and type 5 burst descriptors. DOCSIS 1.x and 2.0 modems use the type 4 burst descriptors; DOCSIS 2.0 modems MUST also use the type 5 burst descriptors. DOCSIS 2.0 modems will use IUCs 9 and 10.

The following rules of operation apply:

- 1) Prior to and during registration a DOCSIS 2.0 TDMA capable modem operating on a channel of type 1 or 2 (refer to clause 11.2.2) MUST calculate its request size based on DOCSIS 1.x IUC parameters. The CMTS MUST make all grants using DOCSIS 1.x IUCs.
- 2) On a type 2 channel, a DOCSIS 2.0 TDMA CM MUST switch to DOCSIS 2.0 TDMA mode after transmission of the Registration Acknowledgement (REG-ACK) message. If the CM receives a Registration Response (REG-RSP) message after transmission of the REG-ACK message, the CM MUST switch back to DOCSIS 1.1 mode before it continues with the registration process (see figure 11-12).
- 3) A CM in DOCSIS 2.0 TDMA mode MUST calculate its request size based on IUC types 9 and 10. The CMTS MUST make grants of IUC types 9 and 10 to that CM after it receives the Registration Acknowledgement message from the CM (see clause 11.2).
- 4) On a type 2 channel, the CM MUST ignore grants with IUCs that are in conflict with its operational mode (e.g. the CM receives a grant with IUC 5 when it is in DOCSIS 2.0 TDMA mode).
- 5) On a type 3 channel, the CMTS MUST use type 5 burst descriptors in order to prevent DOCSIS 1.x modems from attempting to use the channel. All data grants are in IUC types 9 and 10.
- 6) On a type 2 channel, only Advanced PHY Short (IUC 9) and Advanced PHY Long (IUC 10) bursts may be classified as burst descriptor type 5.
- 7) A DOCSIS 1.x modem that does not find appropriate type 4 burst descriptors for long or short data grant intervals MUST consider the UCD and the associated upstream channel, unusable.

G.3.1.2 Interoperability and Performance

This clause addresses the issue of performance impact on the upstream channel when DOCSIS 1.x CMs are provisioned to share the same upstream MAC channel as DOCSIS 2.0 TDMA CMs.

Since the Initial maintenance, Station maintenance, Request and Request/Data IUCs are common to both DOCSIS 2.0 TDMA and DOCSIS 1.x CMs, the overall channel will experience reduced performance compared to a dedicated DOCSIS 2.0 TDMA upstream channel. This is due to broadcast/contention regions not being capable of taking advantage of improved physical layer parameters.

G.3.2 DOCSIS 2.0 S-CDMA Interoperability

G.3.2.1 Mixed mode operation with S-CDMA

In mixed mode operation with both TDMA and S-CDMA, two logically separate upstream channels are allocated by the CMTS, one for TDMA modems and another for DOCSIS 2.0 modems operating in S-CDMA mode. Each channel has its own upstream channel ID and its own UCD. However, these two channels are both allocated the same RF center frequency on the same cable plant segment. The CMTS controls allocation to these two channels in such a way that the channel is shared between the two groups of modems. This can be accomplished by reserving bandwidth through the scheduling of data grants to the NULL SID on all channels other than the channel which is to contain the potential transmit opportunity. Using this method, an upstream channel can support a mixture of differing physical layer DOCSIS modems, with each type capitalizing on their individual strengths. The channel appears as a single physical channel that provides transmission opportunities for both 1.x and DOCSIS 2.0 modems. The mixed-mode configuration of the channel will be transparent to the CMs.

The following rule of operation applies: the CMTS **MUST** use only type 5 burst descriptors on the S-CDMA channel in order to prevent DOCSIS 1.x modems from attempting to use the channel.

G.3.2.2 Interoperability and Performance

This clause addresses the issue of performance impact on the S-CDMA upstream channel when the upstream center frequency is shared with an upstream TDMA channel.

Due to the lack of ability to share the upstream transmit opportunities, the channels will not experience the statistical multiplexing benefits during contention regions across the CMs. Dedicated Initial Maintenance regions will be required on both logical MAC channels slightly reducing the overall performance available. Request and Request/Data regions will also not be capable of being shared although an intelligent CMTS scheduler will be able to reduce most performance impact.

G.3.3 DOCSIS 3.0 Interoperability

A 3.0 CM can initialize on a channel that is described by a Type 35, Type 29 or Type 2 UCD. In the case of a Type 35 UCD, if the CM does not support Selectable Active Code (SAC) Mode 2 and Code Hopping (CH) Mode 2 and the Type 35 UCD has SAC Mode 2 and CH Mode 2 enabled, then the CM **MUST** not use this channel.

Prior to registration, a CM does not operate in Multiple Transmit Channel Mode. Therefore, it follows pre-3.0 DOCSIS rules of requesting as applicable to a Type 1, 2 or 3 channel. Rules regarding Type 2 channels are mentioned in clause G.1.3.

For a Type 4 channel, prior to and during registration a DOCSIS 3.0 cable modem **MUST** calculate its request size in mini-slots based on burst profiles corresponding to IUCs 5 and 6. The CMTS **MUST** make all grants using these burst profiles.

During Registration, if a CM is placed into Multiple Transmit Channel Mode, it transitions to making queue-depth based requests prior to transmission of the REG-ACK message.

If the CM initializes on a Type 4 channel, is not put into Multiple Transmit Channel Mode and DOCSIS 2.0 Mode is enabled, the CM **MUST** begin to calculate its request size based on burst profiles corresponding to IUCs 9 and 10 in the Type 35 UCD beginning after the request for the REG-ACK. The CMTS **MUST** make grants of burst profiles corresponding to IUC 9 and 10 to that CM after it receives the REG-ACK message from the CM (see clause 10.2.6).

G.4 Multicast Support for Interaction with Pre-3.0 DOCSIS Devices

Clause 9.2.2 outlines the CMTS requirements when Multicast DSID Forwarding is enabled on the CMTS. Clause 9.2.2 also outlines the CM requirements when the CMTS sets Multicast DSID Forwarding Capability of '2,' "GMAC-Promiscuous" for the CM.

This clause identifies exceptions or enhancements to the requirements described in clause 9.2.2 for both the CM and CMTS in specific configuration scenarios. These scenarios include:

- "GMAC Explicit DSID Forwarding Mode" in which the CM reports an MDF capability of 1 which is confirmed by the CMTS (clause G.4.2).
- "MDF Mode 0" in which Multicast DSID forwarding is disabled on an MDF-capable CM or a CM is MDF-incapable (clause G.4.3).

G.4.1 Multicast DSID Forwarding (MDF) Capability Exchange

As described in clause 9.2.2, an MDF-capable CM is considered to operate in one of the following three modes of operation based on the value set by the CMTS in REG-RSP or REG-RSP-MP for the Multicast DSID Forwarding (MDF) Capability: "MDF-disabled Mode", "GMAC-Explicit MDF Mode", or "GMAC-Promiscuous MDF mode".

If a CM omits the MDF capability in REG-REQ or REG-REQ-MP (e.g. DOCSIS 2.0 CM), the CMTS omits an MDF encoding in its capability confirmation in REG-RSP or REG-RSP-MP. In addition, a CMTS that does not implement the MDF feature at all (e.g. a CMTS implementing only DOCSIS 2.0 features) sets a value of MDF capability to 0 in REG-RSP or REG-RSP-MP.

The CMTS is allowed to set the value of MDF capability for a CM to 0 in REG-RSP or REG-RSP-MP, irrespective of the value originally reported by the CM in REG-REQ or REG-REQ-MP.

The CMTS is also allowed to set the value of MDF capability to 2 when the CM reports the value of 1 for MDF capability in REG-REQ or REG-REQ-MP. Clause G.4.2.2 provides additional details on this. However, the CMTS is not allowed to set the value of MDF capability to 1 when the CM reports the value of 2 for MDF capability in REG-REQ or REG-REQ-MP.

G.4.2 GMAC-Explicit Multicast DSID Forwarding Mode

GMAC-Explicit MDF Mode means that the CM requires explicit knowledge of the set of multicast Group MAC (GMAC) addresses it is intended to forward. This mode is intended for "Hybrid CMs" that support the ability in hardware to filter downstream unknown GMACs, but do not have the ability in hardware to support filtering of downstream unknown DSID labels. A Hybrid CM is defined as a CM that reports its DOCSIS Version as "DOCSIS 2.0" in its CM Capability Encoding but also separately reports capabilities for selected features of DOCSIS 3.0.

Prior to registration, CMs that report Multicast DSID Forwarding capability as "GMAC Explicit (1)" (clause C.1.3.1.33) are required to forward packets with a destination address of a Well-Known IPv6 MAC address (clause A.1.2) to its IP stack.

A CMTS **MUST** support registration of Hybrid CM that reports a Multicast DSID Forwarding capability as "GMAC Explicit (1)". A Hybrid CM forwards DSID multicast packets according to the forwarding rules associated with the DSID. The CMTS **MUST** by default set this capability with a GMAC Explicit (1) value in the CM Capability Encoding of the REG-RSP or REG-RSP-MP message to the Hybrid CM. A CM to which the CMTS sets the "GMAC Explicit (1)" Multicast DSID Forwarding capability is called a "GMAC-Explicit" Hybrid CM.

When a CMTS adds a DSID on a GMAC-Explicit Hybrid CM, the CMTS MUST include a Multicast Group MAC Address Encoding in the Multicast Encoding, clause C.1.5.4.4.3, for the DSID signaled to that CM. The Multicast Group MAC Address Encoding subtype contains the list of destination Ethernet Group MAC (GMAC) addresses that the CM uses to configure its filter. When the CMTS signals Multicast Group MAC Address Encodings clause C.1.5.4.4.3 to any GMAC-Explicit CM within a DSID Encoding clause C.1.5.3.8, the CMTS MUST NOT label with that DSID any multicast packet that is addressed to GMAC addresses that are NOT signaled in the Multicast Group MAC Address Encoding. This assures that the GMAC-Explicit CM receives all packets labeled with the DSID value.

A Group MAC address becomes a "known Group MAC address" when it is signaled to a Hybrid CM along with an associated DSID. A GMAC-Explicit CM is required to forward downstream multicast packets labeled with a known DSID and with a destination address of a known Group MAC address according to the DSID forwarding rules of clause 9.2.2.3.

For DSID signalling purposes, the GMAC-explicit CM is required to maintain the association between a DSID and a GMAC when they are communicated in the same DSID Encoding (clause C.1.5.3.8). However, this association has no impact on the filtering and forwarding behavior. The DSID and GMAC filters in the GMAC-Explicit CM are independent of each other. Specifically, the GMAC-explicit CM forwards a DSID labeled multicast packet based on the group forwarding attributes of the DSID, as long as both DSID and GMAC are known to the CM, without having to remember the association between two.

This behavior of the GMAC-explicit CM is illustrated by the following example:

- 1) The CMTS signals DSID1 and GMAC1 to the GMAC Explicit CM in the REG-RSP or REG-RSP-MP message along with associated group forwarding attributes such as client MAC address and/or CMIM.
- 2) The CM adds DSID1 in its DSID filter table and GMAC1 in its DMAC filter table.
- 3) The CMTS labels multicast packets with group MAC address of GMAC1 with DSID1 and forwards them to the GMAC-Explicit CM.
- 4) Since both DSID1 and GMAC1 are known to the GMAC-Explicit CM it forwards the packet using group forwarding attributes for DSID1.
- 5) Later on the CMTS signals DSID2 and GMAC2 to the GMAC Explicit CM in the DBC-REQ message along with associated group forwarding attributes such as client MAC address and/or CMIM.
- 6) The CM adds DSID 2 in its DSID filter table and GMAC 2 in its DMAC filter table.
- 7) The CMTS erroneously labels multicast packets with group MAC address of GMAC2 with DSID1 and forwards them to the GMAC-Explicit CM.
- 8) Since the GMAC-explicit CM is not required to maintain the association between the DSID and GMAC for forwarding purposes and both DSID1 and GMAC2 are known to the GMAC-Explicit CM, it forwards the packet using group forwarding attributes associated with DSID1.

G.4.2.1 Example: Forwarding of Multicast Traffic to a Client behind a GMAC-Explicit CM

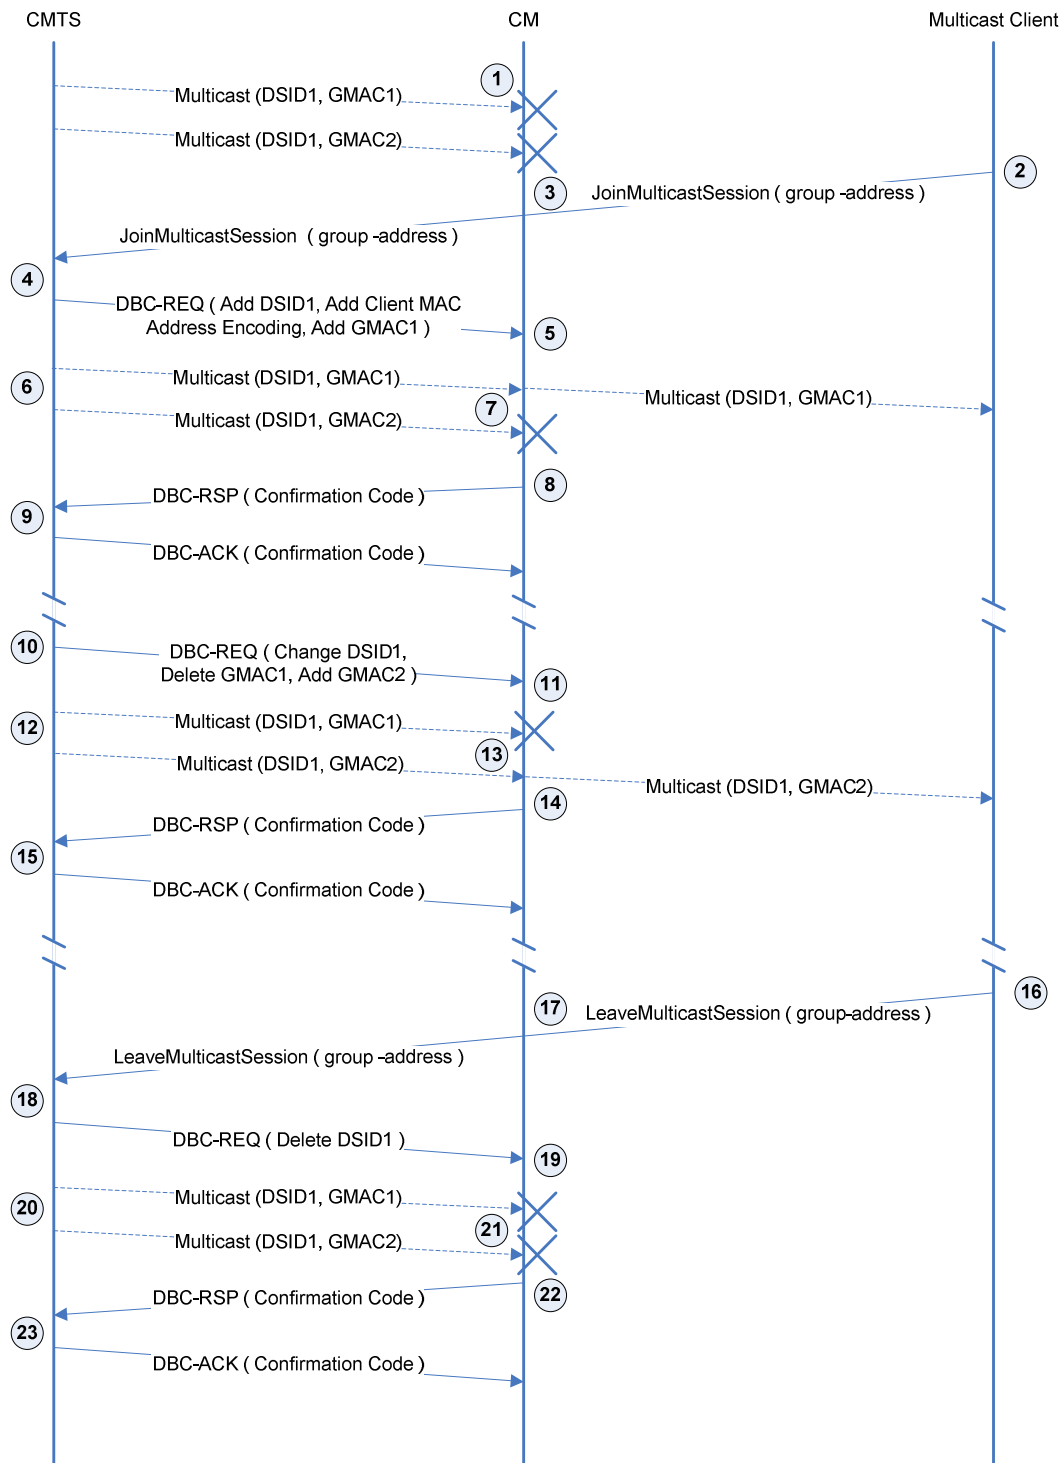


Figure G-1: Multicast Forwarding by a GMAC-Explicit CM

If the CM signaled a preference for GMAC filtering, then the CMTS is required to support the CM request by providing it with the GMAC (Group MAC Addresses) required for filtering incoming multicast packets:

- 1) Multicast packets labeled with DSID1 is not forwarded through the CM to any of the clients, regardless of Group MAC Address (GMAC).
- 2) The Multicast Client sends out a "JoinMulticastSession" when it wants to join an IP Multicast Session.

- 3) The CM forwards the "JoinMulticastSession" upstream to the CMTS like any other data packet without snooping.
- 4) Assuming the CMTS accepts the joiner, the CMTS selects a DSID and sends a DBC-REQ message that includes the DSID, a client MAC address and the GMAC 1 for the IP Multicast Session.

NOTE: The address in the Client MAC address list is the source MAC address in the "JoinMulticastSession" (i.e. the MAC address of the Multicast Client). The CMTS may start sending traffic for that IP Multicast Session labeled with this DSID prior to sending the DBC-REQ message.

- 5) At this point, the CM adds the GMAC 1 to the MAC filter table and the DSID to its DSID filter table. In addition, it associates the client MAC address with this DSID in order to correctly forward multicast packets only to the subscribing Multicast Client.
- 6) When multicast packets arrive at the CMTS, the CMTS labels these packets with the correct DSID and then forwards it downstream.
- 7) In order to minimize the latency, if the CM receives multicast packets it starts forwarding the packets even before it sends DBC-RSP message. When the multicast packet arrives at the CM, the CM only forwards that packet with GMAC1 to the interface on which the Multicast Client is connected (since only this Multicast Client is associated with the DSID signaled to the CM).
- 8) The CM sends DBC-RSP message to the CMTS with appropriate confirmation/error codes.
- 9) The CMTS sends a DBC-ACK message after it successfully receives DBC-RSP message from the CM.
- 10) After some time, for whatever reason the GMAC associated with a DSID needs to be changed. The CMTS sends another DBC-REQ message that Deletes GMAC1 and Adds GMAC2 for that DSID.
- 11) The CM receives the DBC-REQ, removes GMAC1 from its filter table and replaces it with GMAC2.
- 12) When multicast packets destined to GMAC 2 arrive at the CMTS, the CMTS labels these packets with the correct DSID and then forwards the packets downstream.
- 13) In order to minimize the latency, if the CM receives multicast packets it starts forwarding the packets even before it sends DBC-RSP message. When a multicast packet addressed to GMAC2 arrives at the CM, the CM only forwards that packet to the interface on which the Multicast Client is connected (since only this client is associated with the DSID signaled to the CM).
- 14) The CM sends DBC-RSP message to the CMTS with appropriate confirmation/error codes.
- 15) CMTS sends a DBC-ACK message after it successfully receives DBC-RSP message from the CM.
- 16) When the Multicast Client decides to leave the multicast group, it sends a "LeaveMulticastSession".
- 17) The CM forwards the "LeaveMulticastSession" upstream to the CMTS like any other user data packet without snooping.
- 18) Since this is the last Multicast Client behind the CM in the multicast group, the CMTS sends a DBC-REQ to remove the DSID entry in the filter table. Since the GMAC and client MAC address list are associated with a particular DSID, only the DSID needs to be deleted and the GMAC and client MAC address list are deleted along with it.
- 19) The CM receives the DBC-REQ and removes DSID1 and the associated data (GMAC2, Multicast Client 1 MAC address) from the various tables.
- 20) When multicast packets arrive at the CMTS, the CMTS labels it with the correct DSID and then forwards it downstream.
- 21) When the multicast packet arrives at the CM, the CM will not forward that packet.
- 22) The CM sends DBC-RSP message to the CMTS with appropriate confirmation/error codes.
- 23) CMTS sends a DBC-ACK message after it successfully receives DBC-RSP message from the CM.

G.4.2.2 GMAC-Promiscuous Override

A CMTS MAY override the Multicast DSID Forwarding capability of a Hybrid CM from "GMAC-Explicit(1)" to "GMAC-Promiscuous(2)" in the REG-RSP or REG-RSP-MP message to the CM. GMAC Promiscuous forwarding is useful for:

- forwarding a group of IP multicast sessions when any single session is joined;
- forwarding a group of IP multicast sessions to a CPE IP multicast router;
- forwarding all IP multicast sessions with a Layer 2 Virtual Private Network service.

If the CMTS overrides the Multicast DSID Forwarding capability of a Hybrid CM from "GMAC-Explicit(1)" to "GMAC-Promiscuous(2)", the CMTS MUST encrypt all downstream multicast traffic intended to be forwarded by that Hybrid CM with an SAID unique to the DSID label of the multicast traffic. When the CMTS overrides the Multicast DSID Forwarding capability of a Hybrid CM from "GMAC-Explicit(1)" to "GMAC-Promiscuous(2)", the CMTS MUST encrypt all multicast traffic not intended to be forwarded by that Hybrid CM with an SAID unknown to the Hybrid CM. This significantly reduces the performance impact on a CM that is capable of only GMAC-Explicit DSID Forwarding when it is overridden to GMAC-Promiscuous DSID forwarding. Overriding any Hybrid CM to GMAC-Promiscuous DSID forwarding requires the CMTS to encrypt all downstream multicast traffic reaching the Hybrid CM and so makes it mandatory that all CMs in the same MAC domain as the Hybrid CM register with BPI enabled. The CMTS MUST NOT override a Hybrid CM to be in a GMAC Promiscuous (2) mode when any other CM on a MAC domain is not configured to receive encrypted downstream multicast traffic (i.e. if the BPI is not enabled).

G.4.3 MDF Mode 0

A CMTS may implement vendor-specific configuration mechanism to disable MDF on the CMTS globally, on a particular MAC Domain or for particular CMs. The CMTS may return the value 0 for Multicast DSID Forwarding (MDF) capability (clause C.1.3.1.32) in the REG-RSP or REG-RSP-MP to a particular CM to disable MDF for that CM.

Some justifications for a CMTS to disable MDF on some or all CMs capable of supporting it include:

- Globally disabling MDF can reduce the processing and storage requirements on the CMTS in extremely large multicast deployments.
- Existing deployed IPv4 multicast features based on defined DOCSIS 1.1/2.0 IP multicast controls and MIB reporting mechanisms can be maintained while phasing in MDF.

When the CMTS sets the capability of an MDF-capable CM to MDF=0 in the REG-RSP or REG-RSP-MP message, the CM is said to operate with "MDF disabled". CM operation with MDF disabled is specified in clause G.4.3.1.

CMs that either report an MDF capability of zero or do not report an MDF capability (e.g. DOCSIS 1.1/2.0 CMs) are considered to be "MDF-incapable". In this case, the CM forwards multicast per DOCSIS 1.1/2.0 CMs by snooping upstream IGMP v2 joins and forwarding downstream IP multicast packets of the joined sessions. The multicast operation of MDF-incapable CMs is not included in this specification.

G.4.3.1 CMTS Requirements with MDF Mode 0

The following requirements apply to the CMTS when it replicates a multicast session intended to be forwarded through any MDF-disabled or MDF-incapable CM:

- The CMTS MUST omit the Multicast Encoding subtype in any DSID Encoding signaled to an MDF-disabled or MDF-incapable CM (see clause C.1.5.4.4).
- The CMTS MUST NOT replicate a multicast session through an MDF-disabled or MDF-incapable CM with DSID-indexed Payload Header Suppression.
- The CMTS MUST NOT signal to MDF-disabled or MDF-incapable CMs any SAID used for isolating multicast sessions (e.g. bonded multicast) intended to be received by only MDF-enabled CMs.

- The CMTS MUST replicate a multicast session through an MDF-disabled or MDF-incapable CM on only the primary downstream channel of the CM as non-bonded.
- The CMTS MUST transmit a multicast replication through an MDF-disabled or MDF-incapable CM with the Packet PDU MAC Header (Frame Control Type (FC_Type)=00).
- The CMTS MAY omit the DSID label (either by omitting the entire DS-EHDR or by including only 1-byte DS-EHDR) on a multicast replication through an MDF-disabled or MDF-incapable CM.
- The CMTS MAY include a 3-byte DS-EHDR (which includes a DSID label) on the packets of a multicast replication through an MDF-disabled or MDF-incapable CM, even though the CMTS has not signaled the DSID to the MDF-disabled or MDF-incapable CM. This permits the CMTS to use the same replication of a multicast session for MDF-enabled, MDF-disabled and MDF-incapable CMs. The MDF-disabled and MDF-incapable CMs ignore the 3-byte DS-EHDR on multicast packets.
- The CMTS MAY include a 5-byte DS-EHDR on MAC frames of a multicast replication through an MDF-disabled or MDF-incapable CM. This allows the CMTS to use the same replication of a multicast session for MDF-disabled, MDF-incapable and MDF-enabled CMs. In this case, the MDF-enabled CMs recognize the DSID as both a Multicast DSID and a Resequencing DSID. When the CMTS includes a 5-byte DS-EHDR on the MAC frames of a multicast replication through MDF-disabled CMs capable of Multiple Receive Channels, the CMTS MUST signal the DSID to the MDF-disabled CMs as a Resequencing DSID.

NOTE: The CMTS does not signal the DSID as a Multicast DSID to MDF-disabled CMs.

- If the CMTS is configured to disable MDF for all CMs on a MAC Domain, the CMTS MUST transmit pre-registration IPv6 multicast traffic (i.e. intended to be received by the IPv6 host stack of CMs prior to registration) without a DSID label.
- For each CM, the CMTS maintains a supported version of IGMP and MLD. The CMTS MUST maintain the IGMP version as v2 for MDF-disabled and MDF-incapable CMs. The CMTS MUST maintain the MLD version as none for MDF-disabled and MDF-incapable CMs.

The CMTS signals the Security Association of an encrypted multicast session to an MDF-disabled or MDF-incapable CM as defined in [15].

G.4.3.2 CM Requirements with MDF Disabled

The CM operates in the MDF disabled mode when a 3.0 or pre-3.0 CMTS sets the value of MDF capability to 0 in the REG-RSP or REG-RSP-MP. A CMTS regardless of DOCSIS version is required to set the value of unknown CM capability (which in this case is MDF capability) to 0.

The requirements identified in the following clauses are applicable to MDF-disabled CM for backwards compatibility. In accordance with clause 9.1.2.3.2 the CM continues to transparently forward other upstream multicast traffic including, IGMPv3 and MLDv1/v2.

G.4.3.2.1 Requirements for IGMP Management

There are two basic modes of IGMP capability that are applicable to an MDF-disabled CM. The first mode is a passive operation in which the MDF-disabled CM selectively forwards IGMP based upon the known state of multicast session activity on the subscriber side (an example of this is described in annex T. In passive mode, the MDF-disabled CM derives its IGMP timers based on the rules specified in [13]. The second mode is an active operation in which the MDF-disabled CM 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 [57]). A more complete example of an active IGMP device is that of a Multicast Router.

An MDF-disabled CM MUST support IGMPv2 [39]. The CM MUST support IGMP with the cable-specific rules specified in this clause.

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

G.4.3.2.1.1 IGMP Timer Requirements

The following IGMP timer requirements apply only when the MDF-disabled CM is operating in passive IGMP mode:

- The MDF-disabled CM **MUST** be capable of adhering to the timers specified in this clause and not require any specific configuration for the associated multicast timer values.
- The MDF-disabled CM **MAY** provide configuration control that overrides the default values of these timers.
- The MDF-disabled CM **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 $MQI = \text{MAX}(125, MQ_n - MQ_{n-1})$, 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 n th 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.
- The MDF-disabled CM **MUST** support IGMP with the cable-specific rules specified in this clause.
- The MDF-disabled CM **MUST** implement the passive IGMP mode. Additionally, the CM **MAY** implement the active IGMP mode. If it implements the active IGMP mode, the CM **MUST** support a capability to switch between modes.

G.4.3.2.2 Forwarding Requirements for User Joined Multicast

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

- The MDF-disabled CM **MUST** forward traffic for the ALL-SYSTEMS multicast group (224.0.0.1) from its primary downstream to its CPE interface unless administratively prohibited. The CM **MUST** always consider the CPE to be a member of this group. In particular, the MDF-disabled CM **MUST** forward ALL-SYSTEMS Group Queries that pass permit filters on its primary downstream to its CPE interface.
- Upon receiving a Membership Report on its CPE interface, the MDF-disabled CM **MUST** start a random timer between 0 and 3 seconds. During this time period, the MDF-disabled CM **MUST** discard any additional Membership Reports received in its CPE interface for the associated multicast group. If the MDF-disabled CM receives a Membership Report on its HFC interface for the associated multicast group, the MDF-disabled 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 MDF-disabled CMs HFC interface, the MDF-disabled CM **MUST** transmit the Membership Report received on its CPE interface.

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

- The MDF-disabled CM **MUST** implement the Host portion of the IGMP v2 protocol [39] on its RF Interface for CPEs with active groups and **MUST NOT** act as a Querier on its RF Interface.
- The MDF-disabled CM **MUST** act as an IGMPv2 Querier on its CPE interface.
- If the MDF-disabled CM is holding transmission of Membership Reports on its upstream interface for a group active on its CPE interface and receives a Membership Report on its primary downstream, the MDF-disabled CM **MUST** suppress all subsequent Membership Reports for this group until such time as the MDF-disabled CM receives a Membership Query (General or Specific to this Group) on its primary downstream.
- The MDF-disabled CM **MUST** treat Unsolicited Membership Reports (IGMP JOINS) from its CPE interface as a response to a Membership Query received on its primary downstream. Upon receipt of this unsolicited JOIN from its CPE interface, the MDF-disabled CM **MUST** start a random timer according to the Host State Diagram, specified in [39] and use a Query Response Interval of 3 seconds. As specified above, if the MDF-disabled CM receives a Membership Report on its primary downstream for this group during this random time period, it **MUST** suppress transmission of this Join on its upstream RF interface.
- On startup, the MDF-disabled CM **SHOULD** send one or more General Queries on its CPE interface (as described in [39]) in order to quickly and reliably determine membership information for attached CPEs.

- In order to reduce the interruption of multicast service, the MDF-disabled CM SHOULD send a Membership Report for all active multicast groups upon completion of a DCC or DBC operation that involves a downstream channel change.

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

- The MDF-disabled CM MUST NOT forward Membership Queries from its CPE interface to its RF interface.
- The MDF-disabled CM MUST NOT forward Membership Reports or IGMP v2 Leaves received on its primary downstream to its CPE interface.
- The MDF-disabled CM MUST NOT forward multicast traffic from its primary downstream to its CPE interface unless a device on its CPE interface is a member of that IP multicast group.
- The MDF-disabled 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 MDF-disabled CM MUST begin forwarding traffic for the appropriate IP multicast group. The MDF-disabled CM MUST stop forwarding multicast traffic from the primary downstream to the CPE side whenever the MDF-disabled 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 MDF-disabled CM MAY stop forwarding traffic from the primary downstream 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 a IGMP LEAVE message and the appropriate protocol exchange) that there are no CPE devices subscribed to that particular group.
- An MDF-disabled CM MUST discard all downstream traffic (including unicast, multicast and broadcast) with the Isolation Packet PDU MAC Header (Frame Control field (FC_Type) =10).
- An MDF-disabled CM MUST discard downstream multicast traffic when the destination GMAC is unknown. A CM considers any of the following multicast GMAC addresses to be "known":
 - Any Static Multicast MAC Address Encoding configured for the CM.
 - Any DSG MAC address advertised in the DCD message.
 - The multicast MAC address defined via [29], for an IPv4 multicast session joined via an IGMPv2 Join snooped by the CM on a CPE interface.
- An MDF-disabled CM MUST discard all downstream multicast traffic on channels other than its primary downstream channel.
- An MDF-disabled CM MUST ignore a 3 byte DS EHDR on all packets, unicast, multicast and broadcast. An MDF-disabled CM MUST NOT discard packets based on a DSID contained in a 3 byte DS EHDR.
- An MDF-disabled CM MUST discard unicast and broadcast packets with a 5 byte DS EHDR containing an unknown DSID value (even if the MAC address or SAID is known). The CM MUST NOT generate a TEK Invalid (see [15]) due to a key sequence error or report a CRC error in this case.
- An MDF-disabled CM MUST NOT discard downstream multicast packets with a 5-byte DS-EHDR when the DSID is known as a Resequencing DSID to the CM. In this case, whether the CM performs resequencing operation on such multicast packets is vendor specific, because in this mode multicast packets are received only on primary downstream channel.
- An MDF-disabled CM MAY discard downstream multicast packets with a 5-byte DS-EHDR when the DSID is unknown as a Resequencing DSID to the CM.
- When a CM receives a REG-RSP or REG-RSP-MP with MDF capability set to 0 and the REG-RSP or REG-RSP-MP includes Multicast DSID Encodings, the CM MUST reject the REG-RSP or REG-RSP-MP message.
- When a CM with MDF disabled receives Multicast DSID Encodings in a DBC-REQ, the CM MUST reject the DBC-REQ message.

- An MDF-disabled CM **MUST** forward multicast packets addressed to Static Multicast MAC address provided in a configuration file to all CMCI Ports.
- An MDF-disabled CM **MUST** forward IPv4 multicast packets addressed to a group joined via IGMPv2 to the CPE Interface (external CPE interface or internal eSAFE interface) from which the snooped IGMPv2 message was received. An MDF-Disabled CM **SHOULD** forward IPv4 multicast packets addressed to a group joined via IGMPv2 only to the CPE interface from which the snooped IGMPv2 message was received.
- An MDF-disabled CM **MUST NOT** forward to CMCI Ports, multicast packets other than: (1) Addressed to Static Multicast MAC Address and (2) IPv4 multicast address joined via IGMPv2.
- If the CMTS does not include a Pre-Registration DSID in the MDD message, prior to registration the CM **MUST** forward DSID-unlabelled and DSID-labeled multicast packets addressed to Well-Known IPv6 multicast addresses (clause A.1.2) to CM's IPv6 stack.
- If the CMTS does not include a Pre-Registration DSID in the MDD message, prior to registration the CM **MUST** forward DSID-unlabelled and DSID-labeled multicast packets addressed to its Solicited Node MAC addresses to CM's IPv6 stack.
- If the CMTS includes a Pre-Registration DSID in the MDD message, prior to registration the CM forwards multicast packets labeled that Pre-Registration DSID to CM's IPv6 stack (refer to clause 9.2.2). In this case, the CM discards unlabelled multicast packets or multicast packets labeled with other DSIDs.

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

G.4.3.2.3 Forwarding Requirements For Multicast Traffic Associated with IPv6

An MDF-disabled CM **MUST NOT** discard any of the following multicast GMAC addresses used for IPv6 (which are considered to be "known"):

- The well-known IPv6 multicast addresses defined in annex A.
- The Solicited Node multicast MAC addresses corresponding to all IPv6 unicast addresses assigned to the CM IPv6 host stack.
- If IPv6 provisioning of eSAFEs is supported, the Solicited Node multicast MAC addresses corresponding to all IPv6 unicast addresses assigned to the eSAFE IPv6 host stacks.

The following requirements apply to an MDF-disabled CM after the completion of its registration process:

- An MDF-disabled CM **MUST** forward multicast packets (labeled or unlabeled) addressed to Well-Known IPv6 multicast addresses (clause A.1.2) to its IPv6 host stack.
- An MDF-disabled CM **MUST** forward multicast packets (labeled or unlabeled) addressed to the CM's Solicited Node MAC addresses to its IPv6 host stack.
- An MDF-disabled CM **MAY** forward multicast packets (labeled or unlabeled) addressed to Well-Known IPv6 multicast addresses (clause A.1.2) to its eSAFEs.
- An MDF-disabled CM **MAY** forward multicast packets (labeled or unlabeled) addressed to the eSAFEs' Solicited Node MAC addresses to the corresponding interfaces. An MDF-disabled CM does not know the Solicited Node MAC addresses of the CPEs connected to the CMCI Ports as the CM is not expected to learn these addresses by snooping.

Annex H (normative): DHCPv6 Vendor Specific Information Options for DOCSIS 3.0

Please refer to [1], CableLabs DHC Options Registry Specification.

Annex I:
Void

Annex J (normative): DHCPv4 Vendor Identifying Vendor Specific Options for DOCSIS 3.0

Please refer to [1], CableLabs DHC Options Registry Specification.

Annex K (normative): The Data-Over-Cable Spanning Tree Protocol

Clause 9.1 requires the use of the spanning tree protocol on CMs that are intended for commercial use and on bridging CMTSs. This annex describes how the 802.1d spanning tree protocol is adapted to work for data-over-cable systems.

K.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 [16] is employed in most bridged local area networks. This protocol was intended for private LAN use and requires some modification for cable data use.

K.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 [16] 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 K-1 illustrates the general topology.




Figure K-1: Spanning Tree Topology

The task for the public spanning tree protocol, with reference to figure K-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 K-2.




Figure K-2: Spanning Tree Across CMTSs

In figure K-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: 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 needs to be prevented by means external to spanning tree; for example, by using routers.

K.3 Public Spanning Tree Protocol Details

The Data over Cable Spanning Tree algorithm and protocol is identical to that defined in [16], 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 [16]. These BPDUs will be forwarded rather than recalculated by ordinary IEEE 802.1D [16] 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 differentiate further these BPDUs from those used by IEEE 802.1D [16] bridges, in the event that some of those bridges do not correctly identify multicast MAC addresses.
- IEEE 802.1D [16] 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: 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.

K.4 Spanning Tree Parameters and Defaults

Clause 4.10.2 of [16] specifies a number of recommended parameter values. Those values should be used, with the exceptions listed below.

K.4.1 Path Cost

In [16], the following formula is used:

$$\text{Path_Cost} = 1\,000 / \text{Attached_LAN_speed_in_Mb/s}$$

For CMs, this formula is adapted as:

$$\text{Path_Cost} = 1\,000 / (\text{Upstream_modulation_rate} * \text{bits_per_symbol_for_long_data_grant})$$

That is, the modulation type (QPSK or 16 QAM) for the Long Data Grant IUC is multiplied by the raw modulation rate to determine the nominal path cost. Table K-1 provides the derived values.

Table K-1: CM Path Cost

Modulation Rate kHz	Default Path Cost	
	QPSK	16 QAM
160	3 125	1 563
320	1 563	781
640	781	391
1 280	391	195
2 560	195	98

For CMTSs, this formula is:

$$\text{Path_Cost} = 1\,000 / (\text{Downstream_symbol_rate} * \text{bits_per_symbol})$$

K.4.2 Bridge Priority

The Bridge Priority for CMs SHOULD default to 36 864 (0x9000). This is to bias the network so that the root will tend to be at the CMTS. The CMTS SHOULD default to 32 768, as per [16].

NOTE: Both of these recommendations affect only the *default* settings. These parameters, as well as others defined in [16], SHOULD be manageable throughout their entire range through the Bridge MIB [32] or other means.

Annex L (informative): MAC Service Definition

In case of conflict between this clause and any normative clause of this specification, the normative clause takes precedence.

L.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 do not-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.

L.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 clauses 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.

L.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 CMTS Forwarder service. These represent an abstraction of the service provided and do not imply a particular implementation:

- MAC_DATA_INDIVIDUAL.request.
- MAC_DATA_GROUP.request.
- MAC_DATA_INTERNAL.request.
- MAC_DATA.indicate.
- MAC_GRANT_SYNCHRONIZE.indicate.
- MAC_CMTS_MASTER_CLOCK_SYNCHRONIZE.indicate.

L.2.1 MAC_DATA_INDIVIDUAL.request

A CMTS Forwarder issues this primitive to a DOCSIS MAC Domain to forward a packet through an individual CM. This primitive is intended primarily for layer 2 unicast packets, but may also be used to forward an encrypted broadcast or multicast L2PDU through an individual CM.

Parameters:

- CM - the individual CM through which the PDU is intended to be forwarded.
- L2PDU - IEEE 802.3 or [DIX] encoded PDU including all layer two header fields and optional FCS.

Expanded Service Description:

A CMTS Forwarder entity invokes the MAC_DATA_INDIVIDUAL.request primitive of MAC Domain to request the downstream transmission of an L2PDU intended to be forwarded by an individual CM. The mandatory parameters are the L2PDU and an identifier for the individual CM. The L2PDU contains all layer-2 headers, layer-3 headers, data and (optional) layer-2 checksum, but is not considered to contain a DOCSIS Extended Header. This primitive is defined only for Data PDU frame types with Frame Control (FC) values 00 and 10. All MAC Management messages to CMs (with FC=11) are considered to be transmitted by the MAC Domain itself. The MAC Domain is considered to determine and add all DOCSIS Header information.

With this primitive, the packet is classified using the individual Classifier objects instantiated for the individual CM in order to determine the Individual Service Flow with which the MAC Domain schedules downstream transmission for the L2PDU. 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.

This annex does not specify how a CMTS Forwarder component determines the individual CM to which an L2PDU is forwarded. A CMTS forwarder may do so based on the layer 3 IP destination address (if routing), the layer 2 destination MAC address (if bridging) or via some other mechanism (e.g. the encapsulation of the packet when received on an NSI interface, as specified in [8]).

The CMTS Forwarder is considered to deliver a layer 2 PDU to the MAC Domain, so the CMTS Forwarder is responsible for maintaining the IPv4 ARP and IPv6 Neighbor cache table state required to build a Layer 2 PDU from an IP layer 3 datagram. The MAC Domain, however is considered to be responsible for classifying and filtering the L2PDUs based on layer 2 or layer 3 information in the L2PDU.

A CMTS Forwarder is considered responsible for implementing vendor-specific Access Control Lists, while the MAC Domain is responsible for implementing Subscriber Management filtering.

L.2.1.1 Databases

The CMTS MAC Domain is considered to implement a number of databases of objects that persist between packets.

- A database of CABLE_MODEM objects each of which contains all information known in the MAC Domain about the CM. Some attributes of a CABLE_MODEM object CM include:
 - Primary Service Flow ID.

- IsEncrypting -- CM has BPI authorized and active.
- Primary SA -- BPI Security Association for the CM's primary SA.
- A database of INDIVIDUAL_SERVICE_FLOW (ISF) objects indexed by the externally visible Service Flow ID. Some attributes of a downstream individual service flow are:
 - DCS -- Downstream Channel Set on which packets are scheduled.
 - isSequencing -- CMTS is electing to sequence the packets of this ISF.
 - DSID -- DSID label for sequencing the packets of the ISF if the CMTS elects to do so.

NOTE: The CMTS MAY elect to have more than one ISF to the same CM use the same DSID for sequencing.

- A database of INDIVIDUAL_CLASSIFIER_RULE objects associated with an individual CM. Some attributes of an INDIVIDUAL_CLASSIFIER_RULE are:
 - RulePriority -- Priority for matching classifier rule.
 - SFI -- Service Flow ID referenced by the classifier rule.
 - Phsi -- Payload Header Suppression Index that is nonzero if a PHS_RULE is associated with the classifier.
 - Rule Criteria -- criteria for matching an L2PDU submitted for downstream transmission.
 - A database of PHS_RULE objects indexed by CmlId and an 8-bit PHSI. indexed by CM and PHSI.

L.2.1.2 Pseudocode

The following pseudo code describes the intended operation of the MAC_DATA_INDIVIDUAL.request service interface:

```
MAC_DATA_INDIVIDUAL.request (
    CMid,                --internal identifier of a CABLE_MODEM object
    L2PDU)               -- Layer 2 Protocol Data Unit to be txed through the CM
{
  If (the L2PDU matches a downstream subscriber management filter) {
    Discard the packet and return;
  }
}
```

Initialize the DOCSIS Header for the transmitted frame as a non-isolated Data PDU with no extended headers, i.e. with FC_TYPE=00 and FC_PARM=000000.

Attempt to classify the L2PDU with the individual classifier rules of CM.

```
If (L2PDU was matched to an individual classifier){
  Set the transmitting SF to individual SF referenced by the classifier;
  If (the classifier identifies a PHS rule) {;
    Compress the packet using the PHS Rule referenced by the classifier.
  }
} Else {
  Set the transmitting SF to Primary Downstream Service Flow for CM.
}
```

If (the transmitting SF has non-default Traffic Priority) {

```
  Add a 3-byte DS-EDHR to the frame's DOCSIS header, setting the priority bits to the transmitting
  SF's service flow priority;
}
```

Get the Downstream Channel Set (DCS) on which the current frame will be scheduled, as selected by its transmitting SF.

If (the CMTS is sequencing packets from the transmitting SF) {

```
  Get the DSID object for the transmitting ISF;
  Add or increase the DS-EHDR of the transmitted frames DOCSIS Header to use a 5-byte DS-EHDR;
  Set the DS-EDHR's DSID to the transmitting SF's DSID;
```

```

    If (the transmitting ISF is the only ISF for the DSID) {
        Add the next sequence number for the DSID to the DS-EHDR;
        Increment the DSID's sequence number.
    }
}
if (CM is Encrypting) {
    Add a BPI header to the frame using the CM's primary Security Association,
    Encrypt the L2PD using the CM's primary Security Association.
}

```

Enqueue the transmitted MAC frame with the DOCSIS header and L2PDU on the transmitting ISF.

If more than one ISF is using the same DSID, the MAC Domain sets the sequence number of the MAC frame at the time the packet is scheduled to be transmitted, not at the time at which the packet is enqueued for scheduling.

```

} - END MAC_DATA_INDIVIDUAL.request

```

L.2.2 MAC_DATA_GROUP.request

A CMTS Forwarder submits a MAC_DATA_GROUP.request primitive to a MAC Domain in order to forward an L2PDU to an identified group of CMs. This primitive is intended to be used by a CMTS Forwarder primarily to transmit a layer 2 IP multicast packet downstream, but the L2PDU transmitted with this primitive may have a unicast or broadcast destination MAC address. This primitive transmits the packet with a DSID label on the frame.

The primitive has the following parameter variation:

- MAC_DATA_GROUP.request(DCS, L2PDU, DSID).

Where the parameters are:

- DCS - Downstream Channel Set ID to which the L2PDU is replicated.
- L2PDU - IEEE 802.3 or [DIX] encoded protocol data unit starting at the MAC destination address and ending with the last downstream transmitted byte before the FCS.
- DSID - Downstream Service ID that identifies the group of CMs intended to forward the replicated L2PDU.

Prior to invoking this primitive, the CMTS Forwarder initializes the MAC Domain for replicating an IP Multicast Session on a particular DCS of the MAC Domain. The CMTS Forwarder indicates if the IP Multicast Session is encrypted and/or PHS needs to be applied based on the configuration settings. The MAC Domain allocates a Multicast DSID and associates to that Multicast DSID a Security Association and/or DSID-indexed PHS rule. If the DCS is a bonding group, the MAC Domain considers the Multicast DSID as also a Resequencing DSID.

Expanded Service Description:

A CMTS Forwarder entity invokes the MAC_DATA_GROUP.request primitive of MAC Domain to request the downstream transmission of an L2PDU intended to be forwarded by a group of CMs. The L2PDU contains all layer-2 headers, layer-3 headers, data and (optional) layer-2 checksum. It is not considered to contain any DOCSIS Header information; the MAC Domain sub-component adds all DOCSIS Header information to downstream frames.

The MAC_DATA_GROUP.request primitive is intended to describe transmissions to joined IP Multicast groups for which hosts reached through a CM send a Membership Report message in IGMP (for ipv4) or MLD (for ipv6).

The CMTS Forwarder maintains for every (S,G) IP multicast session a set of tuples consisting of MacDomain, DCS and DSID. Each tuple describes how to invoke the MAC_DATA_GROUP.request primitive for replicating the packets of the IP Multicast Session onto a set of DCS.

For transmissions to joined groups, the MAC Domain determines the Group Service Flow (GSF) on which the packet is to be scheduled. The MAC Domain classifies the packet according to a set of Group Classifier Rules (GCRs) associated with the DCS. The GCR refers to the GSF with which the packet is scheduled. The IP Multicast QOS mechanism introduced in DOCSIS 3.0 defines how a Group QOS Table controls the instantiation of GCRs and GSFs when the CMTS forwarder starts replication of an IP multicast session per clause 7.5.8.

This annex does not specify how the CMTS Forwarder component determines how to replicate an IP multicast session, i.e. how the CMTS Forwarder determines the set of (MAC Domain, DCS, DSID) tuples that are used for the parameters of the MAC_DATA_GROUP.request primitive.

The MAC Domain associates with each Multicast DSID the set of CMs to which the Multicast DSID is communicated. The MAC domain associates with each Resequencing DSID a packet sequence number and change count. A Multicast DSID may also be a Resequencing DSID.

The MAC Domain associates a Security Association ID (SAID) with each Multicast DSID used for replicating an encrypted IP Multicast Session.

The following pseudo code describes the intended operation of the MAC_DATA_GROUP.request primitive:

```
MAC_DATA_GROUP.request (
DCSid,  --
L2pdu,
Dsid)
{
Initialize frame's DOCSIS Header with FC_type=00, FC_PARAM= 000000, DS-EHDR field with a length of
3 bytes;
Search the Group Classifier Rules (GCRs) associated with the transmitting DCS for a match to the
L2PDU.
if (matching GCR is found) {
Set the transmitting GSF to the GSF referenced by the matching GCR;
}

Else
{
Set the transmitting GSF to the default GSF for the DCSid
}
Set the Priority field of the DS-EHDR to be transmitted to the Traffic Priority attribute of the
transmitting GSF.
Set the DOCSIS header DSID field to the Dsid parameter of the primitive;
if ( the Multicast DSID is also a Resequencing DSID) {
Set the DS-EHDR to be transmitted to a length of 5;
Set the DS-EHDR's Sequence Change Count to the Resequencing DSID's sequence change count;
Add the Resequencing DSID's packet sequence number to the DS-EHDR;
Increment the Resequencing DSID's packet sequence number;
}
if (the Multicast DSID identifies a DSID-indexed PHS Rule ion) {
Add a PHS Header with PHSI=255 to the DOCSIS Header to be transmitted.
Compress the packet according to the DSID-indexed PHS Rule;
}
if (the Multicast DSID is associated with a Security Association ) {
Add a BPI Header to the DOCSIS Header to be transmitted.
Encrypt the L2PDU with an SA;
}
Schedule the L2PDU with the constructed DOCSIS Header onto the transmitting GSF.
} - MAC_DATA_GROUP.request
```

L.2.3 MAC_DATA_INTERNAL.request

The MAC_DATA_INTERNAL.request primitive represents that CMTS vendors are free to implement any primitive desired for internal data communications between a CMTS Forwarder and the MAC Domain, as long as the subsequent frame transmitted downstream conforms to DOCSIS specifications. In particular, broadcast and multicast packets originated by a CMTS Forwarder, e.g. ARPs, routing advertisements and spanning tree advertisements are not expected to use the defined MAC_DATA_GROUP.request primitive. The CMTS is free to use any CMTS-implemented Group Service Flow (GSF) for CMTS Forwarder initiated multicast packets, but all such packets needs to be accounted for on a GSF.

L.2.4 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.

L.2.5 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.

L.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.

L.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.

L.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 signalling.

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.

L.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.

L.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 this draft of the specification 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.

L.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 signalling.

Parameters:

- ServiceFlowID(s) - unique identifier value(s) for the deleted service flow(s).

L.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.

L.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).

L.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 signalling.

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.

L.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.

L.3.10 MAC_CHANGE_SERVICE_FLOW.indicate

Issued by the DSCIS 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 signalling. DSC signalling 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_IDENTIFIER_S.
- 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.

L.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 is demonstrated by the following scenarios.

L.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.

L.4.2 Reception of PDUs to Upper Layer Service from MAC DATA Service

PDUs are received from the DOCSIS RF.

If a 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 the PDU's 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.

L.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 signalling.

- 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 signalling.

- 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 signalling.

Annex M (informative): Plant Topologies

This clause is informative. In case of conflict between this clause and any normative clause of this specification, the normative clause takes precedence.

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 fibre 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 fibre 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 annex looks at the implementation considerations for these cases. Specifically, the first and last applications are profiled. These examples are meant to illustrate one topology and one implementation of that topology.

M.1 Single Downstream and Single Upstream per Cable Segment

This clause presents an example of a single downstream channel and four upstream channels. In figure M-1, the four upstream channels are on separate fibres or separate wavelengths that each serve 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.

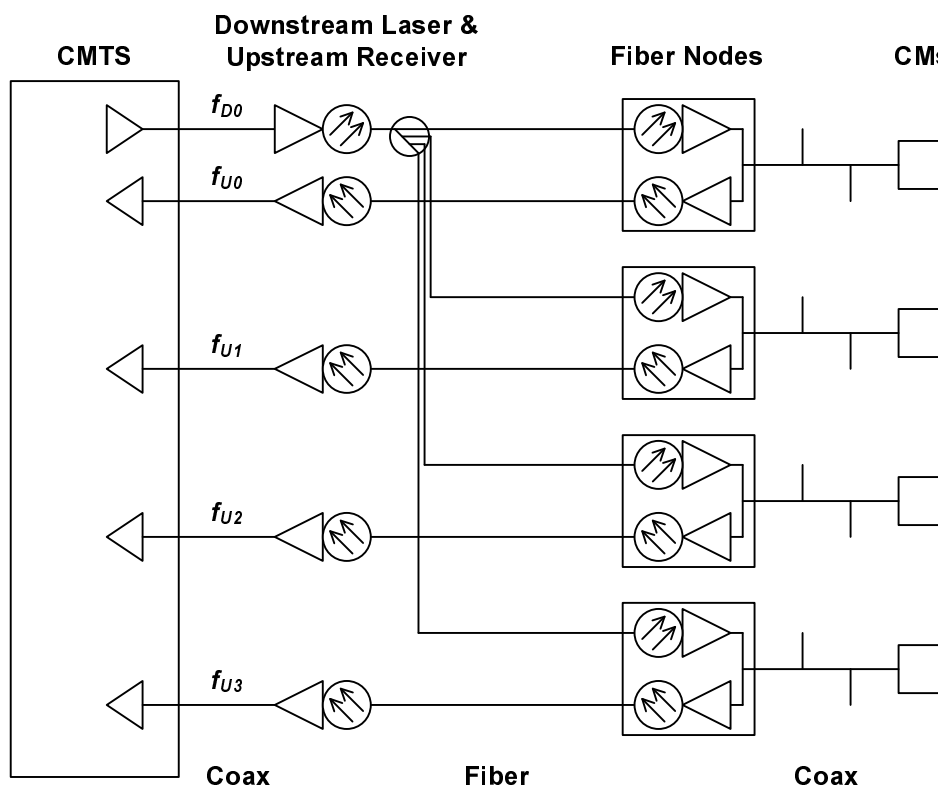


Figure M-1: Single Downstream and Single Upstream Channels per CM

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 fibre 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 then uses 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 then continues 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 fibre branch to which it is connected. If the CM ever has to redo initial ranging, 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 fibre nodes needs to be aligned. If there are multiple logical upstreams sharing the same spectrum on a fibre, then the Initial Maintenance opportunities for each of the logical upstreams needs to align with the Initial Maintenance opportunity of at least one logical upstream with the same center frequency on each fibre node. When the CM chooses a UCD to use and then subsequently uses the MAP for that channel, the CMTS needs to be prepared to receive a Ranging Request at that Initial Maintenance opportunity. Note that only the initialization intervals needs to 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 M-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 fibre node. This could cause initial ranging 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 modulation rates. However, there is a trade-off involved between the time it takes to acquire ranging parameters and flexibility of upstream channel modulation rate. If upstream modulation rates are not the same, the CMTS would be unable to demodulate the Ranging Request if it was transmitted at the wrong modulation rate for the particular upstream receiver of the channel. The result would be that the CM would retry as specified in the [14] 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 modulation rates on different fibre 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 [14] 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 fibre nodes allows the ability to set parameters appropriate for plant conditions on a specific node.

M.2 Multiple Downstreams and Multiple Upstreams per Cable Segment

This clause 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 ranging, 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.

M.2.1 HFC Plant Topologies

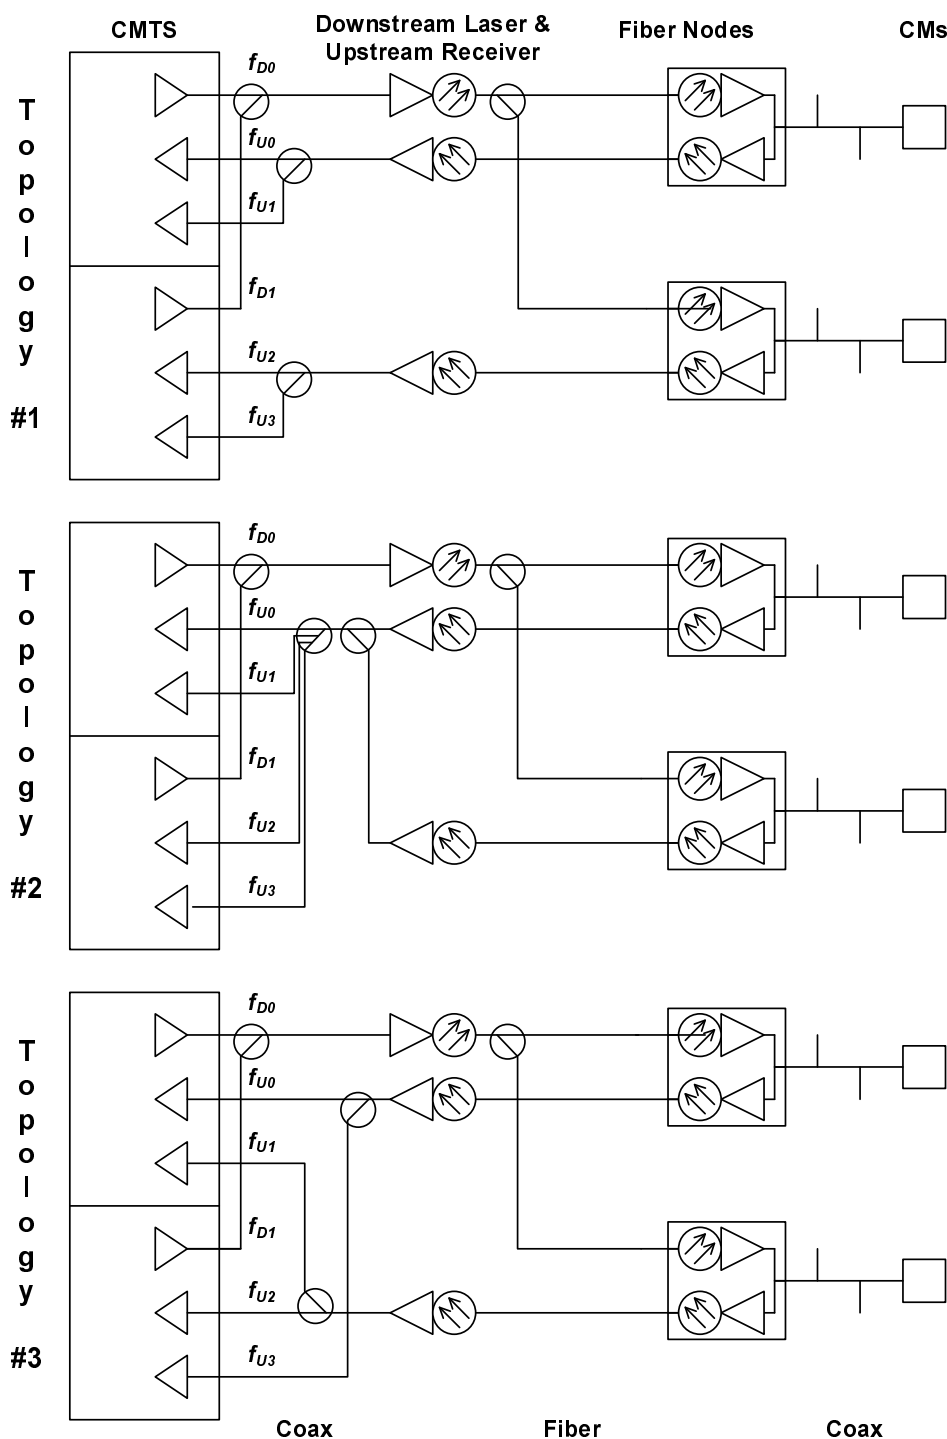


Figure M-2: Bonding Group Example

Suppose two downstream channels are used in conjunction with four upstream channels as shown in figure M-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.

Topology #1 has the return path from each fibre 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 fibre 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 fibre 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.

M.2.2 Normal Operation

Table M-1 lists MAC messages that contain Channel IDs.

Table M-1: MAC Messages with Channel IDs

MAC Message	Downstream 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 needs to 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.

M.2.3 Initial Ranging

When a CM performs initial ranging, 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 needs to be different.
- The constraints specified in clause M.1 apply.

M.2.4 Dynamic Channel Change

With unsynchronized timestamps:

- When a DCC-REQ is given, it needs to 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 needs to 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 fibre 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.

Topologies #1, #2 and #3 will support changes between all upstream and all downstream channels present within the cable segment.

Annex N (informative): DOCSIS Transmission and Contention Resolution

N.1 Multiple Transmit Channel Mode

N.1.1 Introduction

This annex clarifies how the DOCSIS transmission and contention-resolution algorithms work in Multiple Transmit Channel Mode. It contains a few minor simplifications and assumptions, but should be useful to help clarify this area of the specification.

The simplifications include:

- The text does not explicitly discuss packet arrivals while deferring or waiting for pending grants, nor the sizing of piggyback requests.
- The text does not discuss the deferring for a contention request while waiting for grants or grant-pending IEs.
- It shows an example of the operation of the active SID cluster (the SID cluster that the CM can currently use for requests) and an inactive SID cluster (a SID cluster that the CM previously used for requests and for which the CM still has grants pending); the text does not explicitly discuss SID Cluster switching.
- The text does not discuss the possibility of multiple inactive SIDs.

The assumptions include, among others:

- The assumption is made that a Request always fits in any Request region.

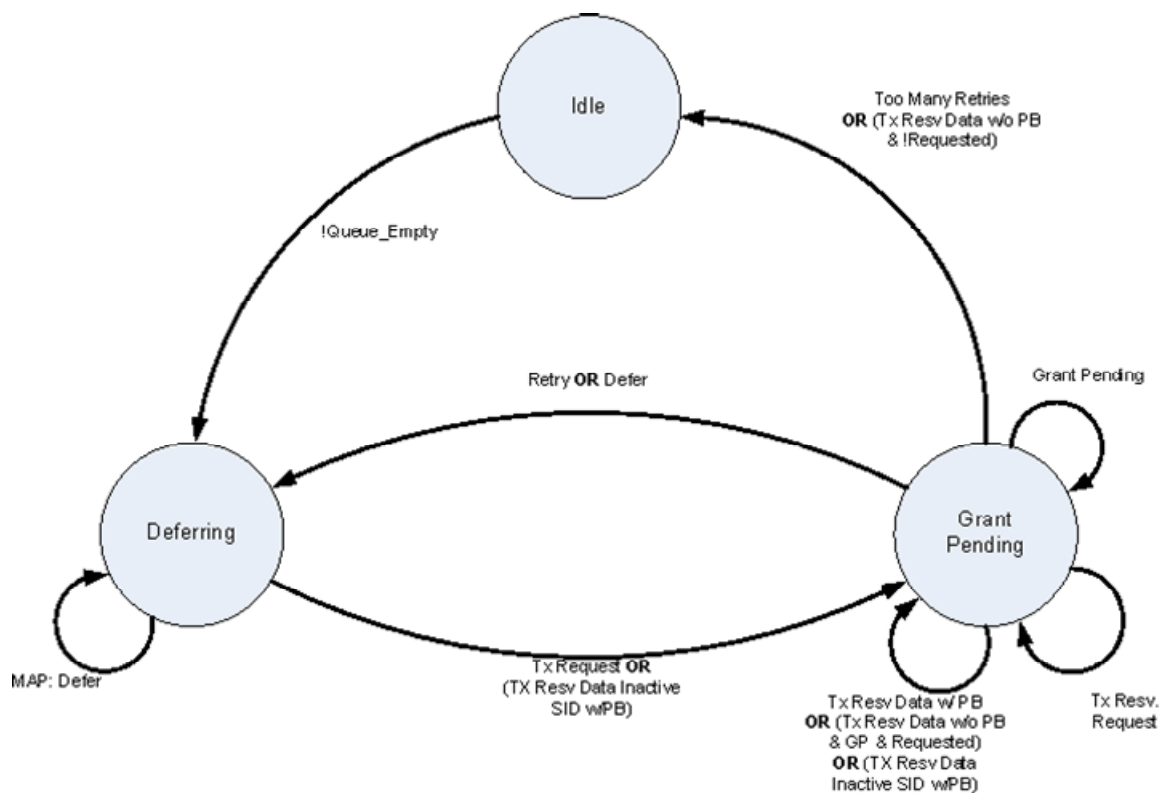


Figure N-1: Transmission and Deference State Transition Diagram (Multiple Transmit Channel Mode)

N.1.2 Variable Definitions

Start [channel i] = Data Backoff Start field from Map "currently in effect" for upstream channel i among the channels associated with the requesting service flow.

End [channel i] = Data Backoff End field from Map "currently in effect" for upstream channel i among the channels associated with the requesting service flow.

Window [channel i] = Current backoff window exponent for upstream channel i among the channels associated with the requesting service flow.

Window_sum = sum of all current backoff windows for all upstream channels in the bonded upstream group.

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 [SID Cluster i] = Saved time of when request was transmitted for SID Cluster i.

Ack_time [SID Cluster i] = Ack Time field from current MAP of upstream channel i.

Piggyback = Flag set whenever a piggyback REQ is available to be sent on the next piggyback opportunity.

Queue_Empty = Flag set whenever the data queue for this service flow does not have un-requested bytes or bytes for which to re-request.

Requested [SID Cluster i] = bytes requested for but not granted yet on SID Cluster i.

Unrequested_bytes = bytes that are in the queue but not requested for yet.

Rerequest_flag = flag indicating if CM failed contention requesting and needs to re-request again for data.

Contention_flag [SID Cluster i] = flag indicating if the SID Cluster i is in contention phase (sent request and waiting for acknowledgement).

Queue_Empty = (unrequested_bytes = 0).

Active_sid = any of the SIDs belonging the SID Cluster that is currently used to send requests.

Inactive_sid = any of the SIDs belonging to the SID Cluster that the CM previously used for requests and for which the CM still has grants pending.

Grant_size_a = number of bytes granted in the current map for a SID belonging to the active SID Cluster.

Grant_size_i = number of bytes granted in the current map for a SID belonging to the inactive SID Cluster.

N = number of upstream channels in the CM's bonded upstream.

Backoff_multiplier = service flow parameter that is the multiplier to the contention request backoff window.

State machine transition definition:

Tx Request = sent request in unicast request opportunity, reserved region or broadcast request opportunity.

Tx Resv. Request = Sent request in a reserved slot.

Tx Resv. Data = received a grant for data.

PB = sent piggyback request in a data grant.

Requested = requested[active_sid] > 0 or requested[inactive_sid] > 0.

GP = grant_pending[active_sid] || grant_pending[inactive_sid].

Defer = look for an opportunity to send request for data.

N.1.3 State Examples

N.1.3.1 Idle - Waiting for a Packet to Transmit

```
Window = 0;
Retries = 0;
Wait for!Queue_Empty; /* Packet available to transmit */
CalcDefer();
go to Deferring
```

N.1.3.2 Grant Pending - Waiting for a Grant

```
Wait for next Map;
Process_map();
utilizeGrant();
stay in state Grant Pending
```

N.1.3.3 Deferring - Determine Proper Transmission Timing and Transmit

```
Wait for next Map;
Process_map();
if (is_my_SID(Grant SID)) /* Unsolicited Grant */
{
    UtilizeGrant();
}
else if (is_my_SID(unicast Request SID) ) /* Unsolicited Unicast Request */
{
    transmit Request in reservation;
    Tx_time[active_sid] = time;
    go to state Grant Pending;
}
else
{
    for (each Request or Request/Data Transmit Opportunity across all MAPS)
        /* request opportunities are counted in time order*/
        {
            if (Defer!= 0)
                Defer = Defer - 1; /* Keep deferring until Defer = 0 */
            else
            {
                transmit Request in contention;
                Tx_time[Active_sid] = time;
                Contention_flag[active_sid] = true;
                go to state Grant Pending;
            }
        }
}
stay in state Deferring
```

N.1.4 Function Examples

N.1.4.1 CalcDefer() - Determine Defer Amount

```
Window_sum = 0;
for (all channels associated with service flow)
{
    if (Window[i] < Start[i])
        Window[i] = Start[i];
    if (Window[i] > End[i])
        Window[i] = End[i];
    Window_sum += 2**Window[i]-1;
}
Defer = Random(floor(backoff_multiplier[]*Window_sum));
```

N.1.4.2 UtilizeGrant() - Determine Best Use of a Grant

```

if (grant_size_a >0) /* CM can send partial or full requested data */
{
    /*reset retries and window*/
    requested[active_sid] -= grant_size_a;
    contention_flag[active_sid] = false;
    if(requested[active_sid] <0)
    {
        Unrequested_bytes += requested[active_sid];
        Requested[inactive_sid] = 0;
        If(unrequested_bytes <0) unrequested_bytes = 0;
    }
}

if (grant_size_i >0) /* CM can send partial or full requested data */
{
    /*reset retries and window*/
    requested[inactive_sid] -= grant_size_i;
    contention_flag[inactive_sid] = false;
    if(requested[inactive_sid] <0)
    {
        Unrequested_bytes += requested[inactive_sid];
        Requested[inactive_sid] = 0;
        If(unrequested_bytes <0) unrequested_bytes = 0;
    }
}

if(unrequested_bytes >0) piggyback = true;

if (requested[active_sid]>0 && !grant_pending[active_sid] && timeout(active_sid))
{
    unrequested_bytes += requested[active_sid];
    if(contention_flag[active_sid] = true)
        rerequest_flag = true;
    piggyback = true;
    requested[active_sid] = 0;
}
if (requested[inactive_sid]>0 && !grant_pending[inactive_sid] && timeout(inactive_sid))
{
    unrequested_bytes += requested[inactive_sid];
    requested[inactive_sid] = 0;
    piggyback = true;
}

for(all grants in this map)
{
    if (active_sid == grant_sid && grant_size_a >0) /* CM can send partial or full requested data */
    {
        transmit max bytes in reservation;
        if(unrequested_bytes >0)
            Tx_time[active_sid] = time;
        unrequested_bytes = 0;
        rerequest_flag = false;
    }

    if (inactive_sid == grant_sid && grant_size_i > 0) /* inactive sid */
    {
        transmit max bytes in reservation;
        if (unrequested_bytes >0)
            Tx_time[active_sid] = time;
        unrequested_bytes = 0;
        rerequest_flag = false;
    }
}

if( piggyback &&(grant_size_a > 0 || grant_size_i > 0)) /* piggyback op was used*/
{
    Piggyback = false;
    Rerequest_flag = 0;
    go to state Grant Pending
}
else if (grant_pending[active_sid] || grant_pending[inactive_sid])
{
    if(grant_pending[active_sid]) contention_flag[active_sid] = false;
    if(grant_pending[inactive_sid]) contention_flag[inactive_sid] = false;
    go to state Grant Pending
}

```



```

    else if(piggyback) /* No grants for this service flow in this map and no grant pendings, no
piggyback op*/
    {
    if(rerequest_flag)
    retry(); /*update number of retries.*/
    else
        go to state Deferring;
    }
else
    go to state Idle

```

N.1.4.3 Retry()

```

Retries = Retries + 1;
if (Retries > 16)
{
discard requested bytes, indicate exception condition
if (QEmpty)
go to state Idle;
}
For (all channels i associated with service flow)
Window[i] = Window[i] + 1;
go to state Deferring;

```

N.1.4.4 Process Map()

```

i = Map.channel_id;
Ack_time[i] = Map.ack_time;
Update grant_pending for active and inactive sid; /* = 0 if no grant-pending IE in current maps from
all channels, otherwise, = 1*/
Grant_size_a = Get the number of bytes granted in this map for active SID
Grant_size_i = Get the number of bytes granted in this map for inactive SID

```

N.1.4.5 timeout (sid)

```

if (min(Ack_time[i], i=0,...,N) > Tx_time[sid])
return true;
else
    return false;

```

N.1.4.6 is_my_SID(sid)

```

If(sid belongs to active SID cluster or inactive SID cluster)
    return true;
return false;

```

N.2 Non-Multiple Transmit Channel Mode

N.2.1 Introduction

This annex clarifies how the DOCSIS transmission and contention-resolution algorithms work when not operating in Multiple Transmit Channel Mode. It contains a few minor simplifications and assumptions, but should be useful to help clarify this area of the specification.

The simplifications include:

- The text does not explicitly discuss packet arrivals while deferring or waiting for pending grants, nor the sizing of piggyback requests.
- The CM always 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 to concatenation, but no attempt is made to address all the subtleties of that situation.

The assumptions include, among others:

- The assumption is made 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 the CMTS.

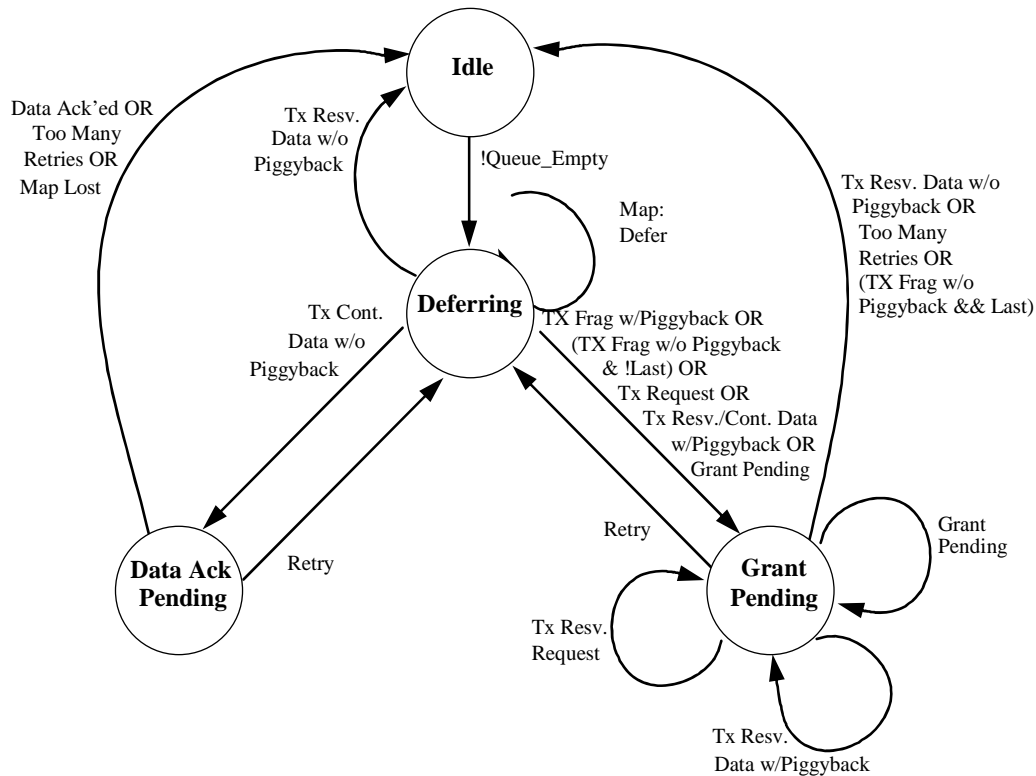


Figure N-2: Transmission and Deference State Transition Diagram

N.2.2 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 and 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.

N.2.3 State Examples

N.2.3.1 Idle - Waiting for a Packet to Transmit

```
Window = 0;
Retries = 0;

Wait for!Queue_Empty; /* Packet available to transmit */
CalcDefer();
go to Deferring
```

N.2.3.2 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;
```

N.2.3.3 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
```

N.2.3.4 Deferring - Determine Proper Transmission Timing and 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

```

N.2.4 Function Examples

N.2.4.1 CalcDefer() - Determine Defer Amount

```

if (Window < Start)
    Window = Start;

if (Window > End)
    Window = End;

Defer = Random[2^Window];

```

N.2.4.2 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)
            piggyback_size = 0
        else
            piggyback_size = pkt_size

        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;
}
```

N.2.4.3 Retry()

```
Retries = Retries + 1;
if (Retries > 16)
{
    discard pkt, indicate exception condition
    go to state Idle;
}

Window = Window + 1;

CalcDefer();

go to state Deferring;
```

Annex O (informative): Unsolicited Grant Services

This annex discusses the intended use of the Unsolicited Grant Service (UGS) and Unsolicited Grant Service with Activity Detection (UGS-AD) and includes specific examples.

O.1 Unsolicited Grant Service (UGS)

O.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 needs to accommodate single or multiple CBR media streams per SID.

For the discussion within this annex, 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.

O.1.2 Configuration Parameters

- Nominal Grant Interval.
- Unsolicited Grant Size.
- Tolerated Grant Jitter.
- Grants per Interval.

Explanations of these parameters and their default values are provided in annex C.

O.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 UGS case of no concatenation and no fragmentation.

O.1.4 Jitter

Figure O-1 shows the relationship between Grant Interval and Tolerated Grant Jitter and shows an example of jitter on subflows.

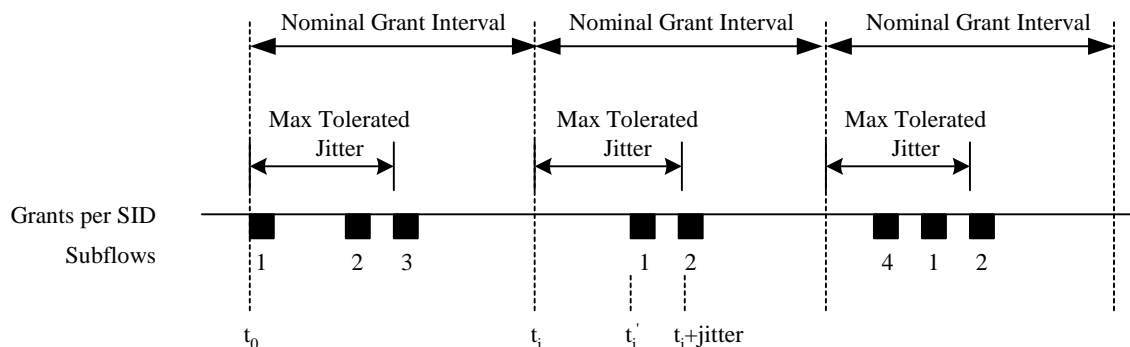


Figure O-1: Example Jitter with Multiple Grants per SID

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 $t_i \leq t'_i \leq t_i + \text{jitter}$.

Figure O-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 to 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.

O.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 the present document. 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 (QI) 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 de-asserts 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.

O.2 Unsolicited Grant Service with Activity Detection (UGS-AD)

O.2.1 Introduction

Unsolicited Grant Service with Activity Detection (UGS-AD) is an Upstream Flow Scheduling Service Type. This clause 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.

For UGS-AD flows, subflows are described as either active or inactive, however the MAC Layer QoS state is still active (i.e. the QoS parameter set is still active).

O.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 annex C.

O.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 needs to 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 needs to be able to restart transmission with either Polled Requests or Unsolicited Grants.

O.2.4 Example




Figure O-2: VAD Start-Up and Stop

Figure O-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.

O.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 packets 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 needs to 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 O-1.

Table O-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 to 1	ms
2	The time until a polled request is received. The worst case time is the Polled Request Interval.	0 to 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 to 15	ms
4	The round trip delay of the HFC plant including the downstream interleaving delay.	1 to 5	ms
Total		6 to 26	ms

This number will vary between CMTS implementations, but reasonable numbers of extra grants to expect from the example above are shown in table O-2.

Table O-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.

O.2.6 Admission Considerations

NOTE: When configuring the CMTS admission control, the following factors needs to 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.

O.3 Multiple Transmit Channel Mode Considerations for Unsolicited Grant Services

In Multiple Transmit Channel Mode, Unsolicited Grant Services can be configured for either segment header-on or segment header-off operation through the Request/Transmission Policy settings. In segment header-off operation, the flow uses only one upstream channel, since there is no way to re-order packets sent on multiple channels. This mode of operation can be more efficient since the overhead of the segment header is not included in each grant.

In Multiple Transmit Channel Mode with segment header-on operation, UGS flows can be assigned to multiple upstream channels. In this scenario, each grant can be placed on a different upstream channel. However, because UGS does not allow for the fragmenting of packets, each grant will be for the full Unsolicited Grant Size.

NOTE: However, that the Unsolicited Grant Size will need to be 8 bytes larger in order to accommodate the segment headers. Also note that even when multiple grants per interval are spread across multiple upstream channels, all of the grants needs to fall within the tolerated jitter for the flow.

Similarly, Extra grants provided to the flow due to assertion of the Queue Indicator or talk spurt bursts can also be scheduled on any of the channels associated with the flow.

Annex P (informative): Error Recovery Examples

In DOCSIS 3.0 the CMTS assumes the majority of the responsibility for recovering from protocol exceptions. In many cases the CM will not try to recover state on its own. Instead, it will wait for the CMTS to direct it on how to recover. This approach allows for CMTS vendor differentiation while maintaining a standard interface between the CM and CMTS. The following examples illustrate how various DOCSIS 3.0 tools can be used to implement this concept.

In the example below not all upstreams were properly ranged before the CM sent the REG-ACK. The CMTS (or an operator or an external program) decides that the best error recovery plan is to instruct the CM to try to range again by re-initializing its MAC.

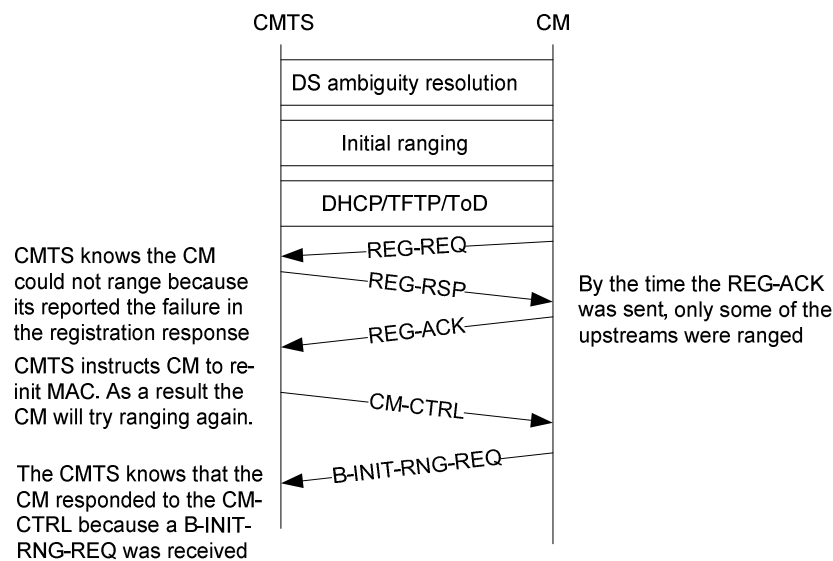


Figure P-1: Example 1 - Modem can not range on all upstreams

In the following example a CM fails to receive MDD messages on one of its non-primary downstream. It reports the error to the CMTS and the CMTS chooses a recovery method. Some legitimate options are:

- 1) Continue to operate in partial mode: A CMTS may choose to take no action. The CM will send 3 CM-STATUS messages and stop. The CM will send a CM-STATUS message with a state "UP" indication as soon as it starts receiving MDD message on the faulty channel again.
- 2) Continue to operate in partial mode (a second option): A CMTS may choose to have the CM operate in partial service mode, but send a DBC for a reduced channel set. In this case the CM will not send a CM-STATUS message when the faulty channel is up again, because its not part of the channel set and therefore the CM is not operating in an errored state.
- 3) A CMTS may force a CM MAC re-initialize by sending a CM-CTRL message (hoping that the reset will correct the error).
- 4) A CMTS may move the CM to a different bonding group which has the same number of channels as the original one. This way service level is not impacted.

Figure P-2 outlines the protocol exchange for option (4).

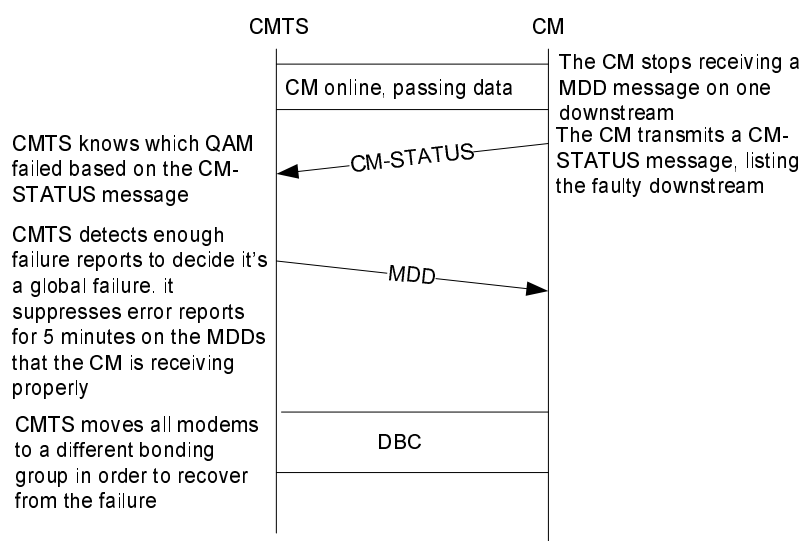


Figure P-2: Example 2 - Option 4

To find a stray modem a CMTS may:

- send DBC with a "null" operation on all DS;
- schedule ranging opportunity and see which upstream responds.

Annex Q (informative): SDL Notation

The SDL (Specification and Description Language) notation used in figure Q-1 is shown in figure VI-1 of ITU-T Recommendation Z.100 [22].

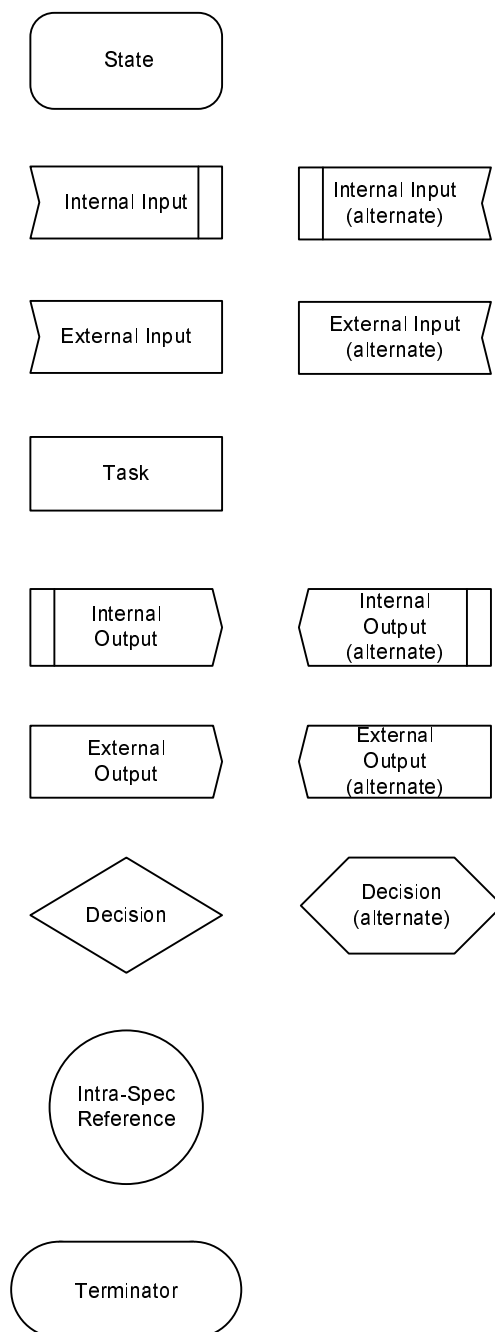


Figure Q-1: Specification and Description Language (SDL) Notation

Annex R (informative):

Notes on Address Configuration in DOCSIS 3.0

DOCSIS 3.0 specifies DHCPv6 as the method of choice to provision IPv6 addresses for CM and bridged devices. [60] defines an alternate mechanism known as stateless address autoconfiguration, where devices build their own IPv6 address by concatenating a prefix learned through router advertisements (RA) and an interface ID derived from the MAC address. Such addresses are usually not registered within the Cable Operator, so their usage is not recommended in DOCSIS 3.0. The simplest way to prevent CM and bridged devices to use stateless address autoconfiguration is to configure router advertisement to not include any prefixes at all.

A CMTS can provide support for enforcing a deployment in which devices attached to the HFC use only DHCPv6 addresses by filtering IPv6 traffic and dropping any IPv6 datagrams whose source address has not been assigned through DHCPv6. Note that this filtering will catch manually assigned address as well as unauthorized SLAAC addresses.

Annex S (informative): IP Multicast Replication Examples

This annex provides examples of some of the key multicast session replication scenarios under mixed CM deployments in the field. It is assumed that the DSID based Multicast Forwarding is enabled on the CMTS in these examples; hence the CMTS always labels multicast packets with a DSID.

When the CMTS replicates a multicast session, it has to make decisions on the following:

- 1) Forwarding the replicated session bonded or non-bonded.
- 2) Downstream Channel Set used for that replication.
- 3) DSID used for that replication.
- 4) Using the Packet PDU MAC Header (FC_Type=00) or the Isolation Packet PDU MAC Header (FC_Type=10).
- 5) If the multicast session is encrypted using a Per-Session SAID to protect the privacy of the multicast content (refer to clause 9.2.6).

In order to make these decisions, the CMTS keeps track of the negotiated value of the Multicast DSID forwarding capability (clause C.1.3.1.30) and the Frame Control Type Forwarding Capability (clause C.1.3.1.31) along with the receive channel set for each registered CM. For a 2.0 or prior DOCSIS CM, the Multicast DSID forwarding capability and the Frame Control Type Forwarding Capability would be 0 and the receive channel set would contain a single downstream channel. When the CMTS has to forward a multicast session through a group of CMs, the CMTS has to ensure that the session is replicated in a way that is consistent with the capability of the group of CMs. Depending upon the negotiated values of CM capabilities; there are three different categories of CMs.

Table S-1: CM Types based on negotiated capabilities

#	CM Type	Multicast DSID Forwarding Capability	Frame Control Type Forwarding Capability
1	CM operating in 2.0/1.1 Mode	0	0
2	Hybrid CM with FC_Type 10 (supporting the Frame Control Type Forwarding Capability)	1	1
3	CM Operating in 3.0 Mode	2	1

If a given session is being replicated more than once for a MAC domain, the CMTS ensures that the CMs do not forward duplicate packets by using Isolation techniques. The CMTS uses the Isolation Packet PDU MAC Header (FC_Type=10) (clause 9.2.2.2.1) to isolate 2.0 or prior DOCSIS CMs from CMs performing Multicast DSID Forwarding. To isolate different sets of CMs performing Multicast DSID Forwarding, the CMTS allocates different DSIDs for each replication and signals only one of those DSIDs to CMs.

S.1 Scenario I: First Multicast Client joiner to a multicast session (Start of a new Multicast Session)

A CMTS may or may not encrypt the multicast session. Some of the reasons for encrypting the multicast session are to prevent forwarding of multicast packets by 1.0 CMs, to prevent duplicate delivery of multicast by the CMs and to protect the privacy of the multicast content.

Under this scenario we need to consider the following three cases based on CM capabilities.

S.1.1 Scenario I - Case 1

Joined Multicast Client is behind a CM incapable of Multicast DSID Forwarding (e.g. 2.0 CM):

- The CM snoops the upstream IGMP messages from a Multicast Client.
- The CM forwards the upstream IGMP messages from a CPE multicast client to the CMTS.
- If BPI is enabled for the CM, the CM sends SA_MAP request to the CMTS as defined in [15].
- If the multicast session is encrypted, then the CMTS sends SA_MAP Reply with the SAID used for the multicast session. If the multicast session is not encrypted then the CMTS sends SA_MAP Reply indicating that there is no SAID for the multicast session.
- CMTS forwards multicast packets non-bonded, labeled with a DSID, FC-Type=00 and encrypted with a Per-Session SAID, if needed.




Figure S-1: Multicast Session Replication for a client behind a 2.0 CM

S.1.2 Scenario I - Case 2

Joined Multicast Client is behind a CM that reports Multicast DSID Forwarding Capability of 1 and Frame Control Type Forwarding Capability of 1 (i.e. Hybrid CM w/ FC-Type 10 Support):

- The CM transparently forwards to the CMTS all upstream IGMP/MLD messages from a Multicast Client.
- The CMTS communicates DSID and GMAC associated with the multicast session to the CM using a DBC Request message. If the multicast session is encrypted, the CMTS also communicates a Per-Session SAID used for encrypting the multicast session using DBC messaging.
- The CMTS may choose to send multicast packets either bonded or non-bonded depending upon Multiple Receive Channel Support capability reported by the CM.
 - Option 1: If the CMTS chooses to send the multicast session non-bonded, it forwards multicast packets labeled with the DSID, FC-Type 00 or 10 and encrypted with the Per-Session SAID for privacy, if needed.
 - Option 2: If the CMTS chooses to send the multicast session as bonded, it forwards multicast packets labeled with the DSID, FC-Type 10 (for isolation from 2.0 or prior DOCSIS CMs) and encrypted with a Per-Session SAID, if needed.

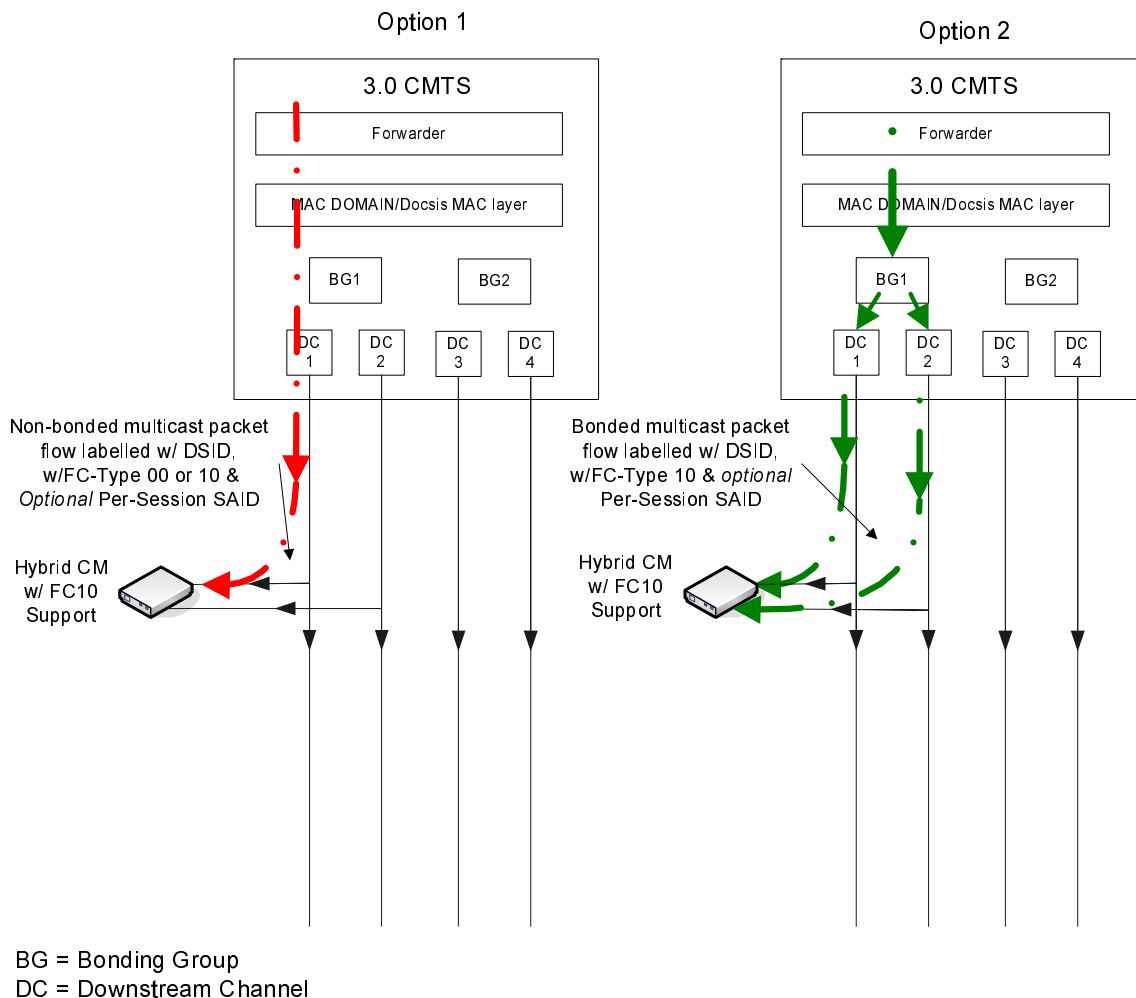


Figure S-2: Multicast Session Replication for a client behind a Hybrid CM capable of FC-Type 10

S.1.3 Scenario I - Case 3

Joined Multicast Client is behind a CM that reports Multicast DSID Forwarding Capability of 2 and Frame Control Type Forwarding Capability of 1 (i.e. 3.0 CM):

- The CM transparently forwards to the CMTS all upstream IGMP/MLD messages from a CPE multicast client.
- The CMTS communicates the DSID associated with the multicast session to the CM using a DBC Request message. If the multicast session is encrypted the CMTS also communicates a Per-Session SAID used for encrypting the multicast session using DBC messaging.
- The CMTS may choose to send multicast packets either bonded or non-bonded depending upon Multiple Receive Channel Support capability reported by the CM.
 - Option 1: If the CMTS chooses to send the multicast session as non-bonded, it forwards multicast packets labeled with the DSID, FC-Type 00 or 10 and encrypted with a Per-Session SAID for privacy, if needed.
 - Option 2: If the CMTS chooses to send the multicast session as bonded, it forwards multicast packets labeled with the DSID, FC-Type 10 (for isolation from 2.0 or prior DOCSIS CMs) and encrypted with a Per-Session SAID for privacy, if needed.

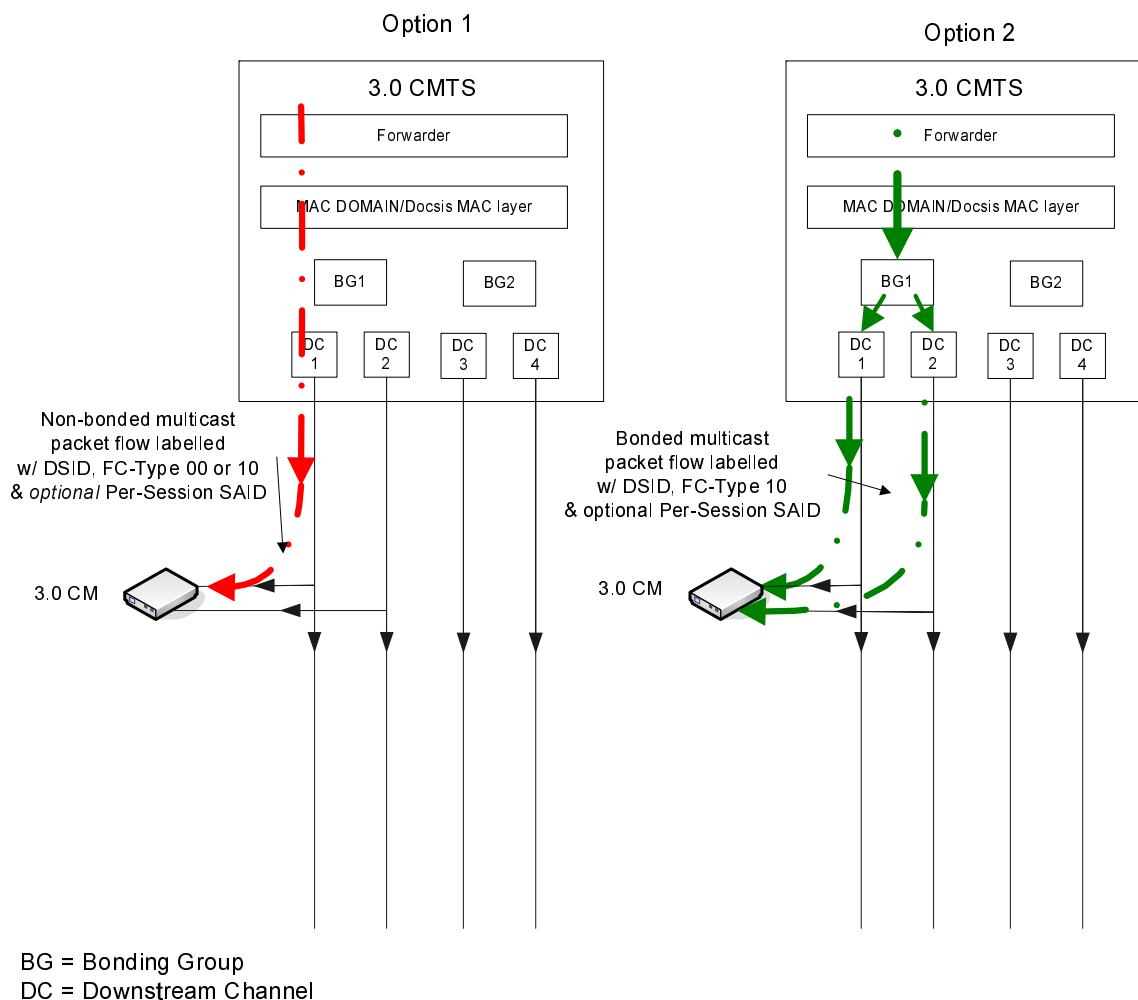


Figure S-3: Multicast Session Replication for a client behind a 3.0 CM

S.2 Scenario II: A Multicast Client joining an existing multicast session that is being forwarded bonded, with FC-Type 10 (Typical 3.0 Multicast Mode of Operation)

At any given moment, the CMTS may be forwarding a multicast session using any one of the techniques outlined under Scenario I, depending upon the capabilities of the CM associated with the first Multicast Client joiner. In addition, a subsequent Multicast Client joiner could be behind a CM that belongs to one of the three different types as outlined above in table S-1. Thus, there can be 9 different combinations under the high-level scenario of subsequent Multicast Clients joining an existing multicast session. However, the following examples cover one specific scenario of a Multicast Client joining an existing multicast session that is being forwarded bonded, labeled with DSID, which is considered as typical DOCSIS 3.0 Multicast Mode of Operation. The CMTS also has the option of forwarding this bonded traffic with either FC-Type 10 or 00. This example covers the case of CMTS forwarding the traffic with FC-Type 10.

A CMTS may or may not encrypt the multicast session. Some of the reasons for encrypting the multicast session are to prevent forwarding of multicast packets by DOCSIS 1.0 CMs, to prevent duplicate delivery of multicast packets by 2.0 or prior DOCSIS CMs and to provide privacy of multicast content.

S.2.1 Scenario II - Case 1

Joined Multicast Client is behind a CM that isn't capable of Multicast DSID Forwarding and can only receive a single downstream channel (e.g. 2.0 CM):

- The CM snoops the upstream IGMP messages from a Multicast Client.
- If BPI is enabled for the CM, the CM sends SA_MAP request to the CMTS as defined in [15].
- If the multicast session is encrypted with Per-Session SAID for privacy, then the CMTS sends SA_MAP Reply with the Per-Session SAID used for the multicast session. If the multicast session is not encrypted then the CMTS sends SA_MAP Reply indicating that there is no SAID for the multicast session.
- Subcase 1: In this case, the CM is tuned to one of the downstream channels on which the Multicast session is currently being forwarded as bonded (either encrypted or unencrypted) and the CMTS chooses to change the multicast session to be forwarded as non-bonded on that downstream channel (may be because there was only one bonding capable CM listening to the multicast session), with FC-Type = 00.

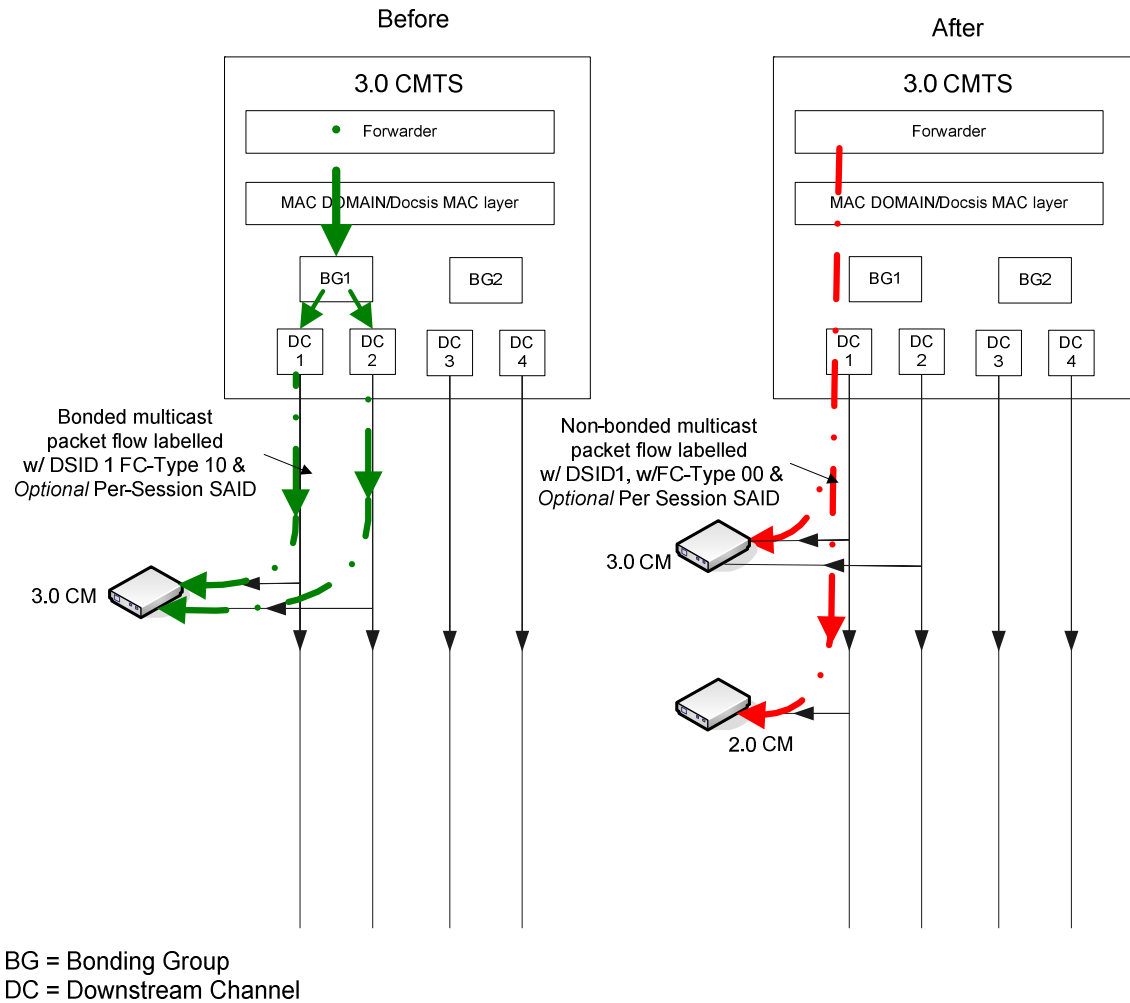


Figure S-4: Multicast Session Replication to clients behind both a 3.0 CM and a 2.0 CM on the same downstream channel (Subcase 1)

- Subcase 2: In this case, the CM is tuned to one of the downstream channels on which the Multicast session is currently being forwarded as bonded (either encrypted or unencrypted) and the CMTS chooses to keep the multicast session as bonded on that downstream channel set, with FC-Type = 10. To accommodate the 2.0 CM, the CMTS needs to add a non-bonded replication with FC-Type = 00. The CMTS uses a DSID not signaled to the 3.0 CMs for the new non-bonded replication to the 2.0CM so that the 3.0 CMs do not forward non-bonded replication. The 2.0 CM will ignore the optional DSID header and forward the packets from the non-bonded replication to the appropriate CPE ports. The 2.0 CM discards the bonded replication since it is sent with FC-Type 10, thus preventing duplicate/partial delivery of multicast packets.




Figure S-5: Bonded and Non-bonded replications of a Multicast Session on an overlapping downstream channel using FC 10 Isolation Technique (Subcase 2)

- Subcase 3: In this case, the CM is not tuned to one of the downstream channels on which the multicast session is currently being forwarded bonded. Hence, the CMTS starts replicating the multicast session on a downstream channel that is received by this new CM as non-bonded with a different DSID (because DSIDs are global to the whole MAC domain), FC-Type 00 and encrypted with the same Per-Session SAID, if needed for privacy.




Figure S-6: Multicast session replications to clients behind both a 3.0 CM and a 2.0 CM on different downstream channel (Subcase 3)

S.2.2 Scenario II - Case 2

Joined Multicast Client is behind a CM that reports Multicast DSID forwarding capability of 1 and Frame Control Type Forwarding Capability of 1 (i.e. Hybrid CM w/ FC-Type 10).

The CM transparently forwards to the CMTS all upstream IGMP/MLD messages from a Multicast Client.

- Subcase 1: In this case, the joining CM can receive the downstream channel set on which the multicast session is being replicated, so the existing multicast session can reach the new joining CM. So CMTS communicates the DSID, Per-Session SAID if the session is encrypted for privacy and GMAC address associated with the multicast session to the newly joined CM using DBC messaging so that the CM can start forwarding the current replication of the multicast session.
- Subcase 2: In this case the new CM cannot receive the downstream channel set on which the multicast session is being replicated, so the CMTS needs to duplicate the multicast session on a different downstream channel set reached by the newly joined CM. The CMTS selects the new DSID for the new replication. CMTS communicates the new DSID, Per-Session SAID, if the session is encrypted for privacy and GMAC address for the new replication of the multicast session to the CM using DBC messaging. The CMTS then starts forwarding the multicast session labeled with the new DSID FC-Type=10 (for isolation from pre-3.0 DOCSIS CMs) and encrypted with the Per-Session SAID for privacy, if needed on the new downstream channel set reached by the newly joined CM.




Figure S-7: Multicast session replication to clients behind both a 3.0 CM and a Hybrid CM w/ FC-Type 10 Support

S.2.3 Scenario II - Case 3

Joined CM reports Multicast DSID Forwarding Capability of 2 and reports that it is capable of FC_Type 10 (3.0 CM):

The CM transparently forwards to the CMTS all upstream IGMP/MLD messages from a CPE multicast client.

- Subcase 1: In this case, the joining CM can receive the downstream channel set on which the multicast session is being replicated, so the existing multicast session can reach the new joining CM. So the CMTS communicates the DSID and per-session SAID, if the session is encrypted for privacy, associated with the multicast session to the newly joined CM using DBC messaging so that the CM can start forwarding the current replication of the multicast session.
- Subcase 2: In this case, the newly joined CM cannot receive the downstream channel set on which the multicast session is being replicated, so the CMTS replicates the multicast session on a different downstream channel set. The CMTS selects the new DSID for this replication. The CMTS communicates the new DSID and per-session SAID, if the session is encrypted for privacy, to the CM using DBC messaging. CMTS then starts forwarding the multicast session labeled with the new DSID, FC-Type=10 and encrypted with the per-session SAID for privacy, if needed, on the new downstream channel set.




Figure S-8: Multicast session replication to clients behind two 3.0 CMs

Annex T (informative): IGMP Example for DOCSIS 2.0 Backwards Compatibility Mode

Clause 9.2.6 defines the requirements for CMTS and CM support of IGMP signalling. This annex provides an example CM passive-mode state machine for maintaining membership of a single multicast group.

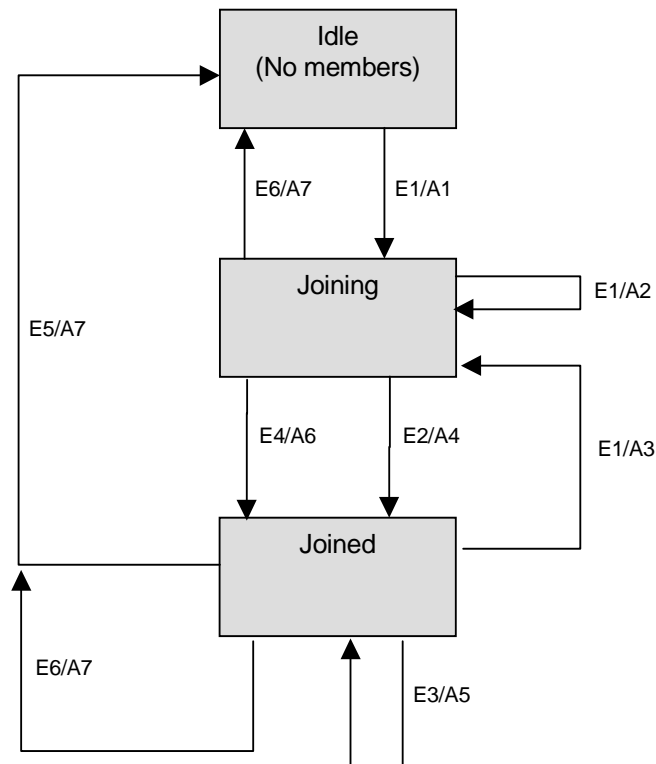


Figure T-1: IGMP Support - CM passive mode

T.1 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.

T.2 Actions

- A1: $MQI = 125$ s; $QRI = 10$ s; Start M1 timer with random value between 0 s and 3 s; start M2 timer = $2 * MQI + QRI$; start TEK machine, if necessary; 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 s and 3 s.
- A4: transmit MR on RF I/f; set I = current time.
- A5: recompute $MQI = \text{MAX}(125, \text{current time} - I)$; set I = current time, forward MQ on CPE i/f.
- A6: cancel M1 timer.
- A7: delete multicast addr from multicast filter.

Annex U (informative): CM Multicast DSID Filtering Summary

The following informational table summarizes the requirements for CMs to drop or forward downstream multicast packets once the CM has completed registration.

Table U-1: CM Post-registration Multicast Filtering Summary

Multicast DSID Forwarding (MDF) Mode	FrameControl	Destination Group MAC	Receive Downstream Channel	DSID Unlabeled	DSID Labeled	
					Unknown as Multicast DSID	Known as Multicast DSID
Pre-DOCSIS 3.0 Multicast (MDF incapable)	FC=00	Known		Forward	Forward	Forward
		Unknown		Drop	Drop	Drop
	FC=10			Drop	Drop	Drop
DOCSIS 3.0 Multicast MDF Mode 0 (MDF disabled)	FC=00	Known	Primary	Forward	Forward	Forward
			Non-Primary	Drop	Drop	Drop
	FC=10	Unknown		Drop	Drop	Drop
				Drop	Drop	Drop
DOCSIS 3.0 Multicast MDF Mode 1 (GMAC Explicit)		Known		Drop	Drop	Forward
		Unknown		Drop	Drop	Drop
DOCSIS 3.0 Multicast MDF Mode 2 (GMAC Promisc.)				Drop	Drop	Forward

The table summarizes the DOCSIS 3.0 requirements for CMs to filter downstream multicast data PDUs under the possible combinations of certain conditions:

- Whether the CM is incapable or capable of MDF Forwarding.
- The Multicast DSID Forwarding (MDF) Mode at which the CMTS confirms an MDF-capable CM to operate.
- The Frame Control value (FC=00) or (FC=10).
- Whether the Destination Group MAC address of the packet is "known" or "unknown". The mechanisms by which a CM learns a GMAC address as known vary depending on the CM's MDF mode.
- Whether the multicast packet is received on the primary or non-primary downstream channel of a CM capable of multiple receive channels.
- Whether the packet is labeled with a DSID or not.
- For DSID-labeled packets, whether the DSID is "known" or "unknown" as a Multicast DSID. Note that when MDF is disabled (MDF mode 0), a DSID is never known as a Multicast DSID.

The table is intended to describe the set of conditions under which the CM is required to filter the packet, denoted by an action of "Drop" in the table. The action denoted by "Forward" means that the CM does not drop the packet for reasons of the conditions in the table. The CM may still drop the packet for other reasons.

For reference, the behavior of CMs operating before DOCSIS 3.0 is also summarized.

A CM with MDF disabled is required to resequence a downstream multicast packet with a 5-byte DS-EHDR only when the DSID of the header identifies a Resequencing Context in the CM. Because an MDF-disabled CMs accepts multicasts only on its primary downstream channel, whether or not the CM actually resequences a packet with a 5-byte DS-EHDR is CM vendor-specific. Whether or not an MDF-disabled CM discards 5-byte DS-EHDR multicast traffic with a DSID unknown as a Resequencing DSID is also vendor specific.

Annex V (informative): Example DHCPv6 Solicit Message Contents

Table V-1: Contents of an example DHCPv6 Solicit message

Option name	Sub-option name	Option code	Contents	Reference
CLIENTID		1	CM DUID	[51], section 22.2 [51], section 9
IA_NA		3		[51], section 22.3
	IAID	(sub-field)	32 bit identifier	
	T1	(sub-field)	0	
	T2	(sub-field)	0	
	IA_NA options	(none)		
VENDOR_CLASS		16	"docsis3.0"	[51], section 22.16
VENDOR_OPTS		17		[51], section 22.17
	ENTERPRISE_NUMBER	(sub-field)	4491	
	ORO	1	Time protocol Time offset TFTP servers Config file name SYSLOG servers	[1]
	TLV5	35	TLV5 attributes as transmitted in MCD	[1], clause C.1.3.1
	DEVICE_ID	36	CM MAC address	[1]

NOTE 1: "sub-field" is a fixed field in the option.
NOTE 2: "none" indicates no suboptions are included.

Annex W (informative): Dynamic Operations Examples

W.1 Dynamic Channel Change Example Operation

W.1.1 Example Signalling

Figure W-1 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 allocate 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 clause 10.2 for more detail.

- 1) An event occurs, such as the CM issuing a DSA-REQ message.
- 2) The CMTS decides that it needs the CM 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 annex C) 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.
- 3) 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.
- 4) 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.
- 5) The CMTS issues a DCC-REQ command to the CM.
- 6) The CM cleans up its queues and state machines as appropriate, sends a DCC-RSP (depart) and changes channels.
- 7) 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.
- 8) If the CM has been instructed to reinitialize, 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 clause 10.2, starting with the recognition of a downstream SYNC message.
- 9) The CM searches for a UCD message unless it has been supplied with a copy.
- 10) The CM waits for a downstream SYNC message unless it has been instructed not to wait for one.
- 11) The CM collects MAP messages unless it already has them available in its cache.
- 12) The CM performs ranging if required by the initialization technique TLV.
- 13) The CM resumes normal data transmission with its new resource assignment.
- 14) The CM sends a DCC-RSP (arrive) message to the CMTS.
- 15) The CMTS responds with a DCC-ACK.
- 16) 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.
- 17) The CM re-issues its DSA-REQ command.
- 18) The CMTS reserves the requested resources and responds with a DSA-RSP.
- 19) The CM finishes with a DSA-ACK.

Note that the flow diagrams for both old and new CMTSs have been carefully crafted to handle many error conditions, such as:

- 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 time out) 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 and 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 s (3 retries of the DCC-REQ).

W.1.2 Example Timing

W.1.2.1 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 TLV's 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 s.
- SYNC interval - 10 ms.
- Unicast ranging interval - 1 s.

The destination CMTS calculates the T15 timer value. The definition of the formula used to determine T15 is shown below. The variables used in calculating T15 are explained in table W-1.

- $T15 = \text{CmJumpTime} + \text{CmtsRxRngReq}.$
- $T15 = 1,3 \text{ s} + (2,04 \text{ s}) = 3,34 \text{ s}.$

Since 3,34 s is less than the minimum value of the T15 timer, the CMTS sets the T15 timer to the minimum value for 4 s.

Table W-1: T15 Calculation Example 1

Variable	Value	Explanation
CmJumpTime	1,3 s	Since the CM did not include the optional jump time TLV, the CMTS will use the default value of 1,3 s.
CmtsRxRngReq	2,04 s $2 * (1 \text{ s}) + 40 \text{ ms}$	Two times the CMTS time period between unicast ranging opportunities plus 20 milliseconds to 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 s.

W.1.2.2 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). It also includes the required UCD substitution TLV and SYNC substitution sub-TLV. 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 s.
- SYNC interval - 10 ms.
- Unicast ranging interval - 5 s.

The destination CMTS starts scheduling the CM immediately after it sends the DCC-REQ. The destination CMTS calculates the T15 timer value. The definition of the formula used to determine T15 is shown below. The variables used in calculating T15 are explained in table W-2.

- $T15 = \text{CmJumpTime} + \text{CmtsRxRngReq}$.
- $T15 = 1,3 \text{ s} + (10,04 \text{ s}) = 11,34 \text{ s}$.

Table W-2: T15 Calculation Example 2

Variable	Value	Explanation
CmJumpTime	1,3 s	Since the CM did not include the optional jump time TLV, the CMTS will use the default value of 1,3 s.
CmtsRxRngReq	10,04 s $2 * (5 \text{ s}) + 40 \text{ ms}$	Two times the CMTS time period between unicast ranging opportunities plus 20 milliseconds to 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, apply the UCD message provided, collect MAP messages on the destination channel without waiting for a downstream SYNC 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 and two ranging intervals ($0 + 2 \text{ s} = 2 \text{ s}$).

W.2 Dynamic Bonding Change Example Operation

W.2.1 Change to Transmit Channel Set and Service Flow SID Cluster Assignments

This is an example in which the CMTS is adding a channel to a Service Flow that requires a modification to the Transmit Channel Set. Figure W-2 describes the sequence of events that happens in the DBC messaging.

In this example, the CM has Service Flows, Service Flow A uses upstream channels 1 and 2 and Service Flow B uses upstream channels 2 and 3. The Transmit Channel Set consists of upstream channels 1, 2 and 3. The CMTS wishes to add upstream channel 4 to the Transmit Channel Set and change Service Flow A to use upstream channels 1, 2 and 4. The CMTS sends the CM a DBC-REQ with TLVs communicating these changes. The CM receives the DBC-REQ message. The CM then enables the transmitter on upstream 4 and adds the new SIDs for upstream 4. After successfully ranging on upstream 4, the CM sends the DBC-RSP to the CMTS indicating that it has made the requested changes and that it is now using upstream 4 for Service Flow A. Once the CMTS receives the DBC-RSP message, it sends the CM a DBC-ACK message and starts allocating grants for Service Flow A over upstream channels 1, 2 and 4.

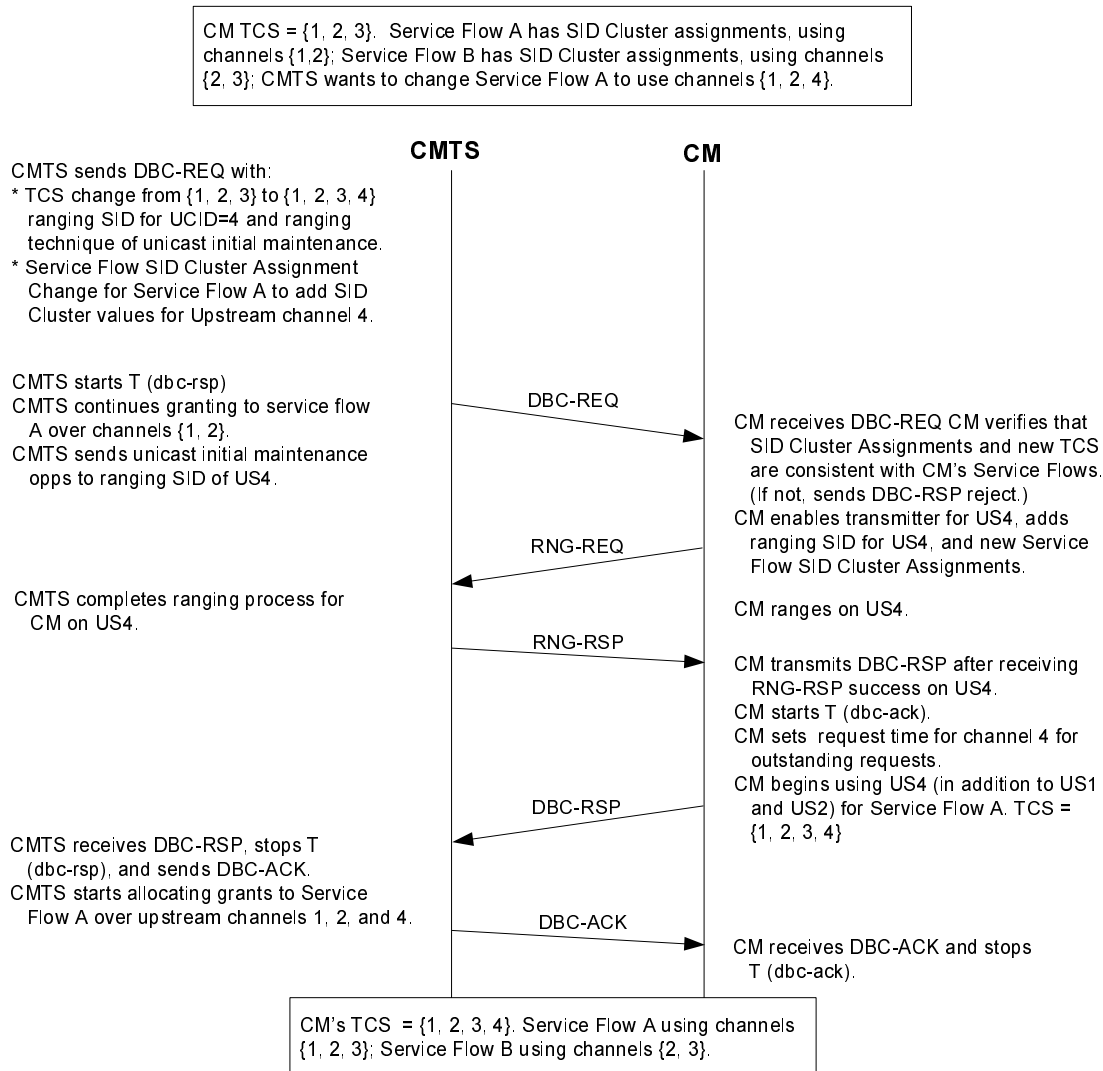


Figure W-2: Adding a Channel to the TCS and making a Service Flow SID Cluster Assignment

W.2.2 Change to Receive Channel Set and Downstream Resequencing Channel List

This is an example in which the CMTS is changing the Downstream Resequencing Channel List of a DSID which requires a modification of the Receive Channel Set. Figure W-3 describes the sequence of events that happen in the DBC transaction.

In this example, the CM has two DSIDs defined, DSID1 and DSID2. Both DSID1 and DSID2 have a Downstream Resequencing Channel List containing downstream channels 1 and 2. The Receive Channel Set consists of downstream channels 1 and 2. The CMTS wishes to add downstream channels 3 and 4 to the Receive Channel Set, move the Downstream Resequencing Channel List of DSID1 from downstream channels 1 and 2 to downstream channels 3 and 4 and expand the Downstream Resequencing Channel List of DSID2 to include downstream channels 3 and 4. The CMTS sends the CM a DBC-REQ with TLVs communicating these changes. The CM receives the DBC-REQ message. The CM stops rapid loss detection of DSID1. The CM then moves the Receive Channel Set to downstream channels 1, 2, 3 and 4, continuing on downstream channels 1 and 2 and acquiring downstream channels 3 and 4. After successfully acquiring downstream channels 3 and 4, the CM sends the DBC-RSP to the CMTS, indicating that it has made the requested changes and is now expecting to receive traffic labeled with DSID1 on downstream channels 1 and 2 and traffic labeled with DSID2 on downstream channels 1, 2, 3 and 4. Once the CMTS receives the DBC-RSP message, the CMTS waits a vendor specific timeout to ensure that the CM receives all data traffic sent prior to the DBC-ACK message, sends the CM a DBC-ACK message, sends traffic associated with DSID1 on downstream channels 3 and 4 and sends traffic associated with DSID2 on downstream channels 1, 2, 3 and 4.

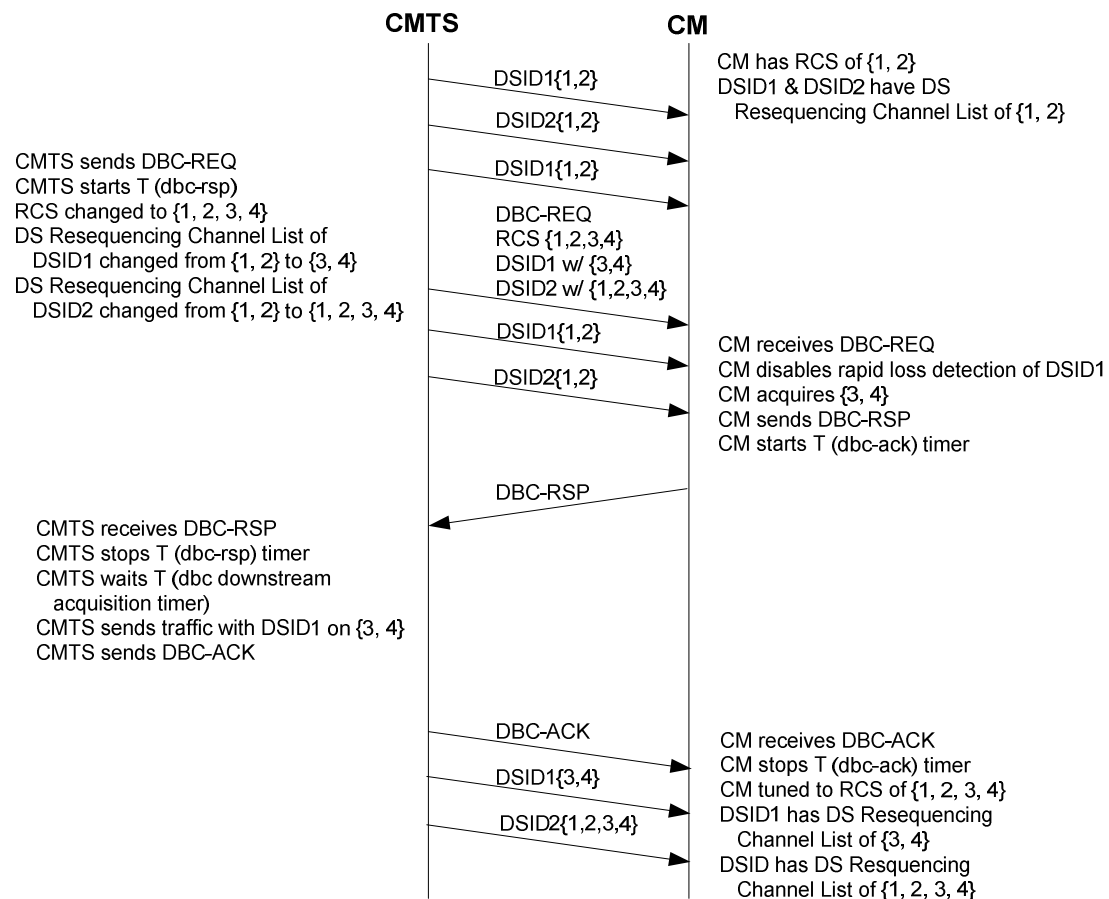


Figure W-3: Changing the RCS and Downstream Resequencing Channel List

W.3 Autonomous Load Balancing Example

Figure W-4 shows an example combining network which illustrates the definition of General Load Balancing Groups and the use of Restricted Load Balancing Groups to resolve topological ambiguities.

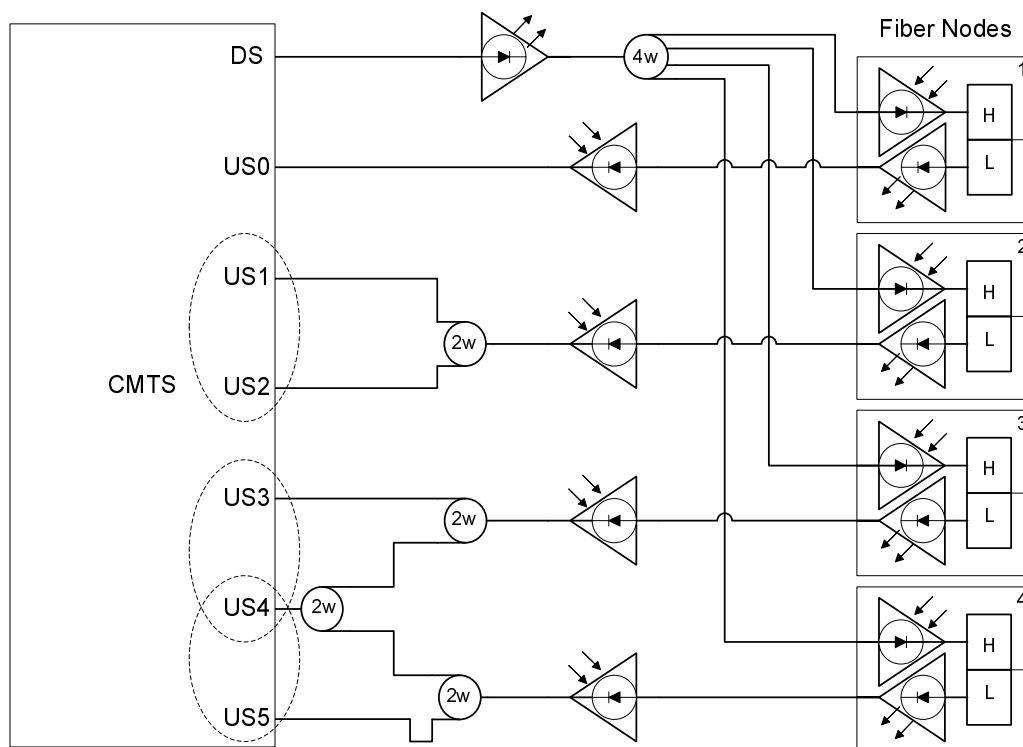


Figure W-4: Example Combining Network 1

In this example, there are six upstream channels (US0 - US5) that are members of a single MAC domain. All six upstream channels are associated with a single downstream channel (DS). The downstream is split over all four fibre nodes, while the six upstreams return from the four nodes via the combining network shown, such that each upstream channel is not physically connected to each fibre node. In particular, fibre node 1 connects to US0 only, fibre node 2 connects to both US1 and US2, fibre node 3 connects to both US3 and US4 and fibre node 4 connects to both US4 and US5.

In this situation, the Load Balancing Groups could be defined as follows:

Load Balancing Group1:

Group ID:	1
Type:	General
Downstream Channels:	DS
Upstream Channels:	US1, US2

Load Balancing Group 2:

Group ID:	2
Type:	Restricted
Downstream Channels:	DS
Upstream Channels:	US3, US4

Load Balancing Group 3:

Group ID:	3
Type:	Restricted
Downstream Channels:	DS
Upstream Channels:	US4, US5

NOTE: A REG-REQ on either upstream channel US1 or US2 uniquely identifies the Load Balancing Group to which a CM can be assigned, hence those two channels form the General Load Balancing Group 1. Upstream channels US3 - US5 have a more complex topology, since US4 is shared across two fibre nodes. To resolve the topological ambiguities that would arise by a REG-REQ received on US4, two Restricted Load Balancing Groups have been defined (Group IDs 2 and 3). In order to be load balanced, each CM that is attached to fibre node 3 would need to be provisioned to be a member of Restricted Load Balancing Group 2, while each CM attached to fibre node 4 would need to be provisioned into Restricted Load Balancing Group 3. If a CM were to register on one of these channels without having been provisioned into the appropriate Restricted Load Balancing Group, the CMTS would not associate the CM with any Load Balancing Group (which results in the CM not being load balanced).

Also, note that US0 is not a member of any Load Balancing Group. CMs which register on that upstream channel will not be load balanced to another channel.

Figure W-5 shows a second example, in which two MAC domains are shared across two fibre nodes in a complex combining network. In this example, a pair of upstream channels (one from each MAC domain) are set aside for a particular customer group (e.g. business customers), a Restricted Load Balancing Group is formed to allow load balancing for those customers.

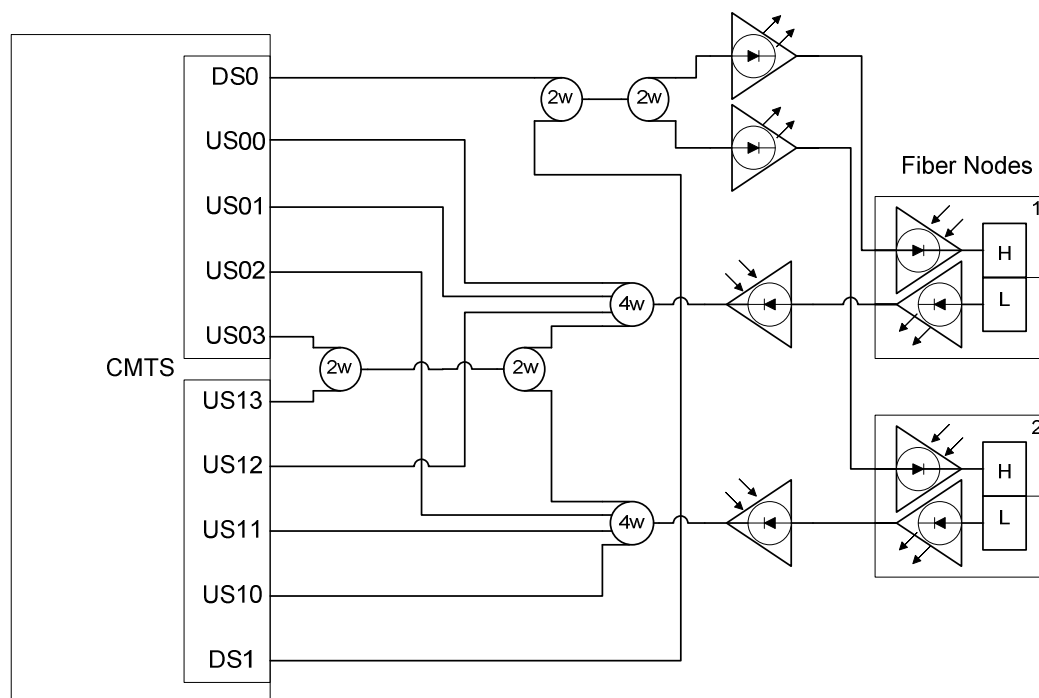


Figure W-5: Example Combining Network 2

Load Balancing Group 1:

Group ID:	1
Type:	General
Downstream Channels:	DS0, DS1
Upstream Channels:	US00, US01, US12
Subgroup:	DS0, US00, US01

Load Balancing Group 2:

Group ID:	2
Type:	General
Downstream Channels:	DS0, DS1
Upstream Channels:	US10, US11, US02
Subgroup:	DS1, US10, US11

Load Balancing Group 3:

Group ID:	3
Type:	Restricted
Downstream Channels:	DS0, DS1
Upstream Channels:	US03, US13

Annex X (informative): Bibliography

- Internet Assigned Numbers Authority, Internet Multicast Addresses.

NOTE: Available at <http://www.iana.org/assignments/multicast-addresses/multicast-addresses.xml> and <http://www.iana.org/assignments/ipv6-multicast-addresses/ipv6-multicast-addresses.xml>.

- IETF RFC 2462 (December 1998): "IPv6 Stateless Address Autoconfiguration", S. Thomson, T. Narten.
- IETF RFC 3219 (January 2002): "Telephony Routing over IP (TRIP)", J. Rosenberg, H. Salama, M. Squire.
- IETF RFC 3495 (March 2003): "Dynamic Host Configuration Protocol (DHCP) Option for CableLabs Client Configuration", B. Beser, P. Duffy.
- IETF RFC 4604 (August 2006): "Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast", H. Holbrook, B. Cain, B. Haberman.
- ETSI ES 201 488-3 (V1.2.2): "Access and Terminals (AT); Data Over Cable Systems; Part 3: Baseline Privacy Plus Interface Specification".

History

Document history		
V1.1.0	April 2011	Public Enquiry PE 20110818: 2011-04-20 to 2011-08-18
V1.1.0	September 2011	Vote V 20111112: 2011-09-13 to 2011-11-14
V1.1.1	November 2011	Publication