



***Society of Cable
Telecommunications
Engineers***

**ENGINEERING COMMITTEE
Data Standards Subcommittee**

AMERICAN NATIONAL STANDARD

ANSI/SCTE 137-2 2017

**Modular Head End Architecture
Part 2: M-CMTS Downstream External PHY Interface**

NOTICE

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

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

Attention is called to the possibility that implementation of this document may require the use of subject matter covered by patent rights. By publication of this document, no position is taken with respect to the existence or validity of any patent rights in connection therewith. If a patent holder has filed a statement of willingness to grant a license under these rights on reasonable and nondiscriminatory terms and conditions to applicants desiring to obtain such a license, then details may be obtained from the standards developer. SCTE shall not be responsible for identifying patents for which a license may be required or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention.

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

All Rights Reserved

© Society of Cable Telecommunications Engineers, Inc. 2017
140 Philips Road
Exton, PA 19341

DOCSIS® and M-CMTS™ are registered trademarks of Cable Television Laboratories, Inc., and used in this document with permission.

Contents

MODULAR HEAD END ARCHITECTURE.....	1
1 SCOPE	8
1.1 SCOPE AND PURPOSE	8
1.2 REQUIREMENTS AND CONVENTIONS	8
2 REFERENCES.....	10
2.1 NORMATIVE REFERENCES.....	10
2.2 INFORMATIVE REFERENCES	11
2.3 REFERENCE ACQUISITION	11
3 TERMS AND DEFINITIONS.....	12
4 ABBREVIATIONS AND ACRONYMS	15
5 TECHNICAL OVERVIEW	18
5.1 SYSTEM ARCHITECTURE	18
5.1.1 <i>Reference Architecture</i>	18
5.1.2 <i>DEPI Operation</i>	19
5.1.3 <i>EQAM Operation</i>	20
5.2 BONDING SERVICES MODEL	21
5.3 MULTIPLE SERVICES MODEL	21
6 DEPI ARCHITECTURE.....	23
6.1 DEPI DATA PATH	23
6.1.1 <i>DOCSIS D-MPT Data Path</i>	24
6.1.2 <i>PSP Data Path</i>	24
6.1.3 <i>DOCSIS SYNC Message</i>	24
6.1.3.1 <i>SYNC Message Format</i>	24
6.1.3.2 <i>Correction and Insertion</i>	25
6.1.4 <i>Latency and Skew Requirements</i>	26
6.1.4.1 <i>Latency</i>	26
6.1.4.2 <i>Skew</i>	26
6.2 NETWORKING CONSIDERATIONS.....	26
6.2.1 <i>Per Hop Behavior Usage</i>	26
TABLE 6–1 - PHBS AND RECOMMENDED DSCP VALUES.....	27
6.2.2 <i>DiffServ Code Point Usage</i>	27
6.2.3 <i>Packet Sequencing</i>	27
6.2.4 <i>Network MTU</i>	27
6.3 SYSTEM TIMING CONSIDERATIONS	28
7 DEPI CONTROL PLANE	29
7.1 TOPOLOGY	29
7.2 ADDRESSING	29
7.3 CONTROL MESSAGE FORMAT	31
7.3.1 <i>Control Message with a UDP Header</i>	32
7.3.2 <i>Control Message without a UDP Header</i>	33
7.3.3 <i>Common Headers for Control and Data Messages</i>	33
7.3.3.1 <i>Ethernet 802.3 Header</i>	33
7.3.3.2 <i>Ethernet 802.1Q Header</i>	34
7.3.3.3 <i>IPv4 Header</i>	34
7.3.3.4 <i>UDP Header</i>	34
7.3.3.5 <i>CRC</i>	34

7.3.4	<i>Specific Headers for Control Messages</i>	34
7.3.4.1	L2TPv3 Control Header.....	34
7.3.4.2	Attribute Value Pairs (AVP)	35
7.4	SIGNALING.....	36
TABLE 7-1 - DEPI CONTROL MESSAGES		36
7.4.1	<i>Control Connection Signaling</i>	37
7.4.1.1	Control Connection Setup	37
7.4.1.2	Control Connection Teardown	38
7.4.1.3	Control Connection Keep-Alive.....	39
7.4.2	<i>Session Signaling</i>	39
7.4.2.1	Session Setup	39
7.4.2.2	Session Teardown	40
7.4.2.3	Session Updates	41
7.4.3	<i>Required and Optional AVPs</i>	41
TABLE 7-2 - DEPI MANDATORY AND OPTIONAL AVPS		42
7.5	AVP DEFINITIONS.....	42
7.5.1	<i>Conventional L2TPv3 AVPs</i>	42
TABLE 7-3 - DEPI SUPPORTED L2TPV3 AVPS		43
7.5.1.1	Message Type (All Messages)	43
7.5.1.2	Result Code (StopCCN, CDN).....	44
7.5.1.3	Host Name (SCCRQ, SCCRP).....	44
7.5.1.4	Vendor Name (SCCRQ, SCCRP)	44
7.5.1.5	Serial Number (ICRQ, OCRQ)	44
7.5.1.6	Router ID (SCCRQ, SCCRP).....	45
7.5.1.7	Assigned Control Connection ID (SCCRQ, SCCRP, StopCCN)	45
7.5.1.8	Pseudowire Capabilities List (SCCRQ, SCCRP)	45
TABLE 7-4 - PSEUDOWIRE TYPES		45
7.5.1.9	Local Session ID (ICRQ, ICRP, ICCN, CDN, SLI).....	46
7.5.1.10	Remote Session ID (ICRQ, ICRP, ICCN, CDN, SLI)	46
7.5.1.11	Remote End ID (ICRQ)	46
7.5.1.12	Pseudowire Type (ICRQ).....	46
7.5.1.13	L2-Specific Sublayer (ICRQ, ICRP, ICCN)	47
TABLE 7-5 - L2-SPECIFIC SUBLAYER TYPES		47
7.5.1.14	Data Sequencing (ICRP).....	47
7.5.1.15	Circuit Status (ICRQ, ICRP, ICCN, SLI)	47
7.5.2	<i>DEPI Specific AVPs</i>	48
TABLE 7-6 - DEPI DEFINED GENERAL SESSION APVS		48
7.5.2.1	DEPI Result Code (CDN, SLI)	48
7.5.2.2	DEPI Resource Allocation Request (ICRQ).....	49
7.5.2.3	DEPI Resource Allocation Reply (ICRP)	49
7.5.2.4	DEPI Local MTU (ICRQ).....	50
7.5.2.5	Downstream QAM Channel DOCSIS SYNC Control (ICRQ, SLI)	50
7.5.2.6	EQAM Capabilities (ICRP).....	51
7.5.2.7	DEPI Remote MTU (ICRP)	51
7.5.2.8	Local UDP Port (ICRQ).....	52
7.5.3	<i>QAM Channel PHY AVPs</i>	52
TABLE 7-7 - DEPI DEFINED QAM CHANNEL PHY AVPS		52
7.5.3.1	Downstream QAM Channel TSID Group (ICRP).....	53
7.5.3.2	Downstream QAM Channel Frequency (ICRP, ICCN, SLI).....	53
7.5.3.3	Downstream QAM Channel Power (ICRP, ICCN, SLI)	54
7.5.3.4	Downstream QAM Channel Modulation (ICRP, ICCN, SLI).....	54
7.5.3.5	Downstream QAM Channel J.83 Annex (ICRP, ICCN, SLI)	54

7.5.3.6	Downstream QAM Channel Symbol Rate (ICRP, ICCN, SLI).....	55
7.5.3.7	Downstream QAM Channel Interleaver Depth (ICRP, ICCN, SLI)	55
7.5.3.8	Downstream QAM RF Block Muting (ICRP, ICCN)	56
8	DEPI FORWARDING PLANE	57
8.1	L2TPv3 TRANSPORT PACKET FORMAT.....	57
8.1.1	<i>Data Message with a UDP Header</i>	58
8.1.2	<i>Data Message without a UDP Header</i>	59
8.1.3	<i>Specific Headers for Data Messages</i>	59
8.1.3.1	L2TPv3 Data Header.....	59
8.1.3.2	L2TPv3 DEPI Sub-Layer Header.....	60
8.1.3.3	DEPI Payload.....	60
8.2	DOCSIS MPT SUB-LAYER HEADER.....	60
8.3	PSP SUB-LAYER HEADER	61
8.4	DEPI LATENCY MEASUREMENT (DLM) SUB-LAYER HEADER.....	62
8.5	M-CMTS CORE OUTPUT RATE.....	63
ANNEX A	DEPI MTU	64
A.1	L2TPv3 LOWER LAYER PAYLOAD SIZE.....	64
A.2	MAXIMUM FRAME SIZE FOR DEPI	64
TABLE A-1	- MTU OF DEPI	64
A.3	PATH MTU DISCOVERY	65
ANNEX B	PARAMETERS AND CONSTANTS.....	66
TABLE B-1	- PARAMETERS AND CONSTANTS.....	66
ANNEX C	DOCS-IF-M-CMTS-MIB (NORMATIVE).....	67
ANNEX D	FORMAT AND CONTENT FOR EVENT, SYSLOG, AND SNMP NOTIFICATION (NORMATIVE)	124
D.1	EVENT DEPI PROCESS DEFINITIONS.....	124
D.2	DEPI EVENTS.....	124
TABLE D-1	- DEPI EVENTS	124
APPENDIX I	DEPI AND DOCSIS SYSTEM PERFORMANCE.....	127
I.1	INTRODUCTION	127
I.2	ROUND-TRIP TIME AND PERFORMANCE.....	127
I.3	ELEMENTS OF ROUND-TRIP TIME	127
TABLE I-1	- DOCSIS REQUEST-GRANT ROUND TRIP WORKSHEET	129
I.4	CIN CHARACTERISTICS.....	129
I.5	QUEUING DELAYS IN NETWORK ELEMENTS.....	130
I.6	TRAFFIC PRIORITIZATION AND NETWORK DELAYS	131
I.7	QUEUE PERSISTENCE IN A DEPI FLOW	131
I.8	PSP MODE	133
APPENDIX II	EARLY ADOPTION AND EVOLVING USE OF EQAM DEVICES.....	134
II.1	EQAM DEVELOPMENT: CATEGORY A (NO DTI).....	134
II.2	EQAM DEVELOPMENT: CATEGORY B (WITH DTI)	134
II.3	POSSIBLE M-CMTS FEATURE PHASING	135
II.4	OPTIONAL UDP LAYER.....	135

List of Figures

FIGURE 5-1 - MODULAR CMTS REFERENCE ARCHITECTURE.....	18
FIGURE 5-2 - EQAM BLOCK DIAGRAM.....	20
FIGURE 5-3 - BONDING SERVICES MODEL.....	21
FIGURE 5-4 - MULTI-SERVICE MODE.....	21
FIGURE 6-1 - DOWNSTREAM EQAM BLOCK DIAGRAM.....	23
FIGURE 6-2 - FORMAT OF A DOCSIS SYNC MAC MESSAGE.....	25
FIGURE 7-1 - L2TP TOPOLOGY FOR MODULAR CMTS.....	29
FIGURE 7-2 - DEPI ADDRESSING HIERARCHY.....	30
FIGURE 7-3 - DEPI ADDRESSING HIERARCHY.....	31
FIGURE 7-4 - DEPI CONTROL PACKET WITH UDP.....	32
FIGURE 7-5 - DEPI CONTROL PACKET WITHOUT UDP.....	33
FIGURE 7-6 - DEPI CONTROL CONNECTION SETUP.....	37
FIGURE 7-7 - DEPI CONTROL CONNECTION TEARDOWN.....	38
FIGURE 7-8 - DEPI KEEP-ALIVE.....	39
FIGURE 7-9 - DEPI SESSION SETUP.....	39
FIGURE 7-10 - DEPI SESSION TEARDOWN.....	40
FIGURE 7-11 - DEPI SESSION UPDATES.....	41
FIGURE 7-12 - MESSAGE TYPE AVP.....	43
FIGURE 7-13 - RESULT CODE AVP.....	44
FIGURE 7-14 - HOST NAME AVP.....	44
FIGURE 7-15 - VENDOR NAME AVP.....	44
FIGURE 7-16 - SERIAL NUMBER AVP.....	44
FIGURE 7-17 - ROUTER ID AVP.....	45
FIGURE 7-18 - CONTROL CONNECTION ID AVP.....	45
FIGURE 7-19 - PSEUDOWIRE CAPABILITIES LIST AVP.....	45
FIGURE 7-20 - LOCAL SESSION ID AVP.....	46
FIGURE 7-21 - REMOTE SESSION ID AVP.....	46
FIGURE 7-22 - REMOTE END ID AVP.....	46
FIGURE 7-23 - PSEUDOWIRE TYPE AVP.....	46
FIGURE 7-24 - L2-SPECIFIC SUBLAYER AVP.....	47
FIGURE 7-25 - DATA SEQUENCING AVP.....	47
FIGURE 7-26 - CIRCUIT STATUS AVP.....	47
FIGURE 7-27 - DEPI RESULT AND ERROR CODE AVP.....	48
FIGURE 7-28 - DEPI RESOURCE ALLOCATION REQUEST AVP.....	49
FIGURE 7-29 - DEPI RESOURCE ALLOCATION REPLY AVP.....	49
FIGURE 7-30 - DEPI LOCAL MTU AVP.....	50
FIGURE 7-31 - DOCSIS SYNC AVP.....	50
FIGURE 7-32 - EQAM CAPABILITIES AVP.....	51
FIGURE 7-33 - DEPI REMOTE MTU MAX PAYLOAD AVP.....	51
FIGURE 7-34 - LOCAL UDP PORT AVP.....	52
FIGURE 7-35 - TSID GROUP AVP.....	53
FIGURE 7-36 - FREQUENCY AVP.....	53
FIGURE 7-37 - POWER AVP.....	54
FIGURE 7-38 - MODULATION AVP.....	54
FIGURE 7-39 - J.83 ANNEX AVP.....	54
FIGURE 7-40 - SYMBOL RATE AVP.....	55
FIGURE 7-41 - INTERLEAVER DEPTH AVP.....	55
FIGURE 7-42 - RF MUTE AVP.....	56
FIGURE 8-1 - L2TPv3 DATA PACKET OUTER ENCAPSULATION WITH UDP.....	58
FIGURE 8-2 - L2TPv3 DATA PACKET OUTER ENCAPSULATION WITHOUT UDP.....	59
FIGURE 8-3 - DOCSIS MPT SUB-LAYER HEADER AND PAYLOAD.....	60
FIGURE 8-4 - DEPI PSP SUB-LAYER HEADER AND PAYLOAD.....	61
FIGURE 8-5 - DLM SUB-LAYER HEADER.....	62

List of Tables

TABLE 6–1 - PHBs AND RECOMMENDED DSCP VALUES.....	27
TABLE 7–1 - DEPI CONTROL MESSAGES	36
TABLE 7–2 - DEPI MANDATORY AND OPTIONAL AVPs.....	42
TABLE 7–3 - DEPI SUPPORTED L2TPv3 AVPs.....	43
TABLE 7–4 - PSEUDOWIRE TYPES	45
TABLE 7–5 - L2-SPECIFIC SUBLAYER TYPES.....	47
TABLE 7–6 - DEPI DEFINED GENERAL SESSION APVs.....	48
TABLE 7–7 - DEPI DEFINED QAM CHANNEL PHY AVPs.....	52
TABLE A–1 - MTU OF DEPI	64
TABLE B–1 - PARAMETERS AND CONSTANTS.....	66
TABLE D–1 - DEPI EVENTS	124
TABLE I–1 - DOCSIS REQUEST-GRANT ROUND TRIP WORKSHEET.....	129

1 SCOPE

1.1 Scope and Purpose

NOTE: This document is identical to SCTE 137-2 2010 except for informative components which may have been updated such as the title page, NOTICE text, headers and footers. No normative changes have been made to this document.

This specification is part of the DOCSIS® family of specifications, and in particular, is part of a series of specifications that define a Modular Cable Modem Termination System (M-CMTS™) architecture for head-end components that comply with DOCSIS. This specification was developed for the benefit of the cable industry, and includes contributions by operators and vendors from North America, Europe, and other regions.

The DOCSIS Specifications [RFI2.0] define the requirements for the two fundamental components that comprise a high-speed data-over-cable system: the cable modem (CM) and the cable modem termination system (CMTS). The M-CMTS architecture was designed as an extension to the DOCSIS Specifications to allow for flexibility and independent scaling of certain CMTS functions, and to allow operators to more efficiently use available network resources.

One of the key elements of the M-CMTS architecture is the separation of the downstream physical layer QAM modulation and up-conversion functions from the CMTS, and the placement of that functionality into an "Edge-QAM" (EQAM) device. This separation allows for the development of EQAM products that support both video and DOCSIS, which in turn allows operators to use the same network resources to support multiple types of services such as data, voice, and video.

This document defines an interface known as the Downstream External PHY Interface (DEPI) and associated protocol requirements for the transport of downstream user data between the "M-CMTS Core" and the EQAM. It describes the characteristics of the DEPI interface, provides requirements that must be met by the M-CMTS Core and the EQAM, and also describes various aspects of technical issues that are involved in the implementation and deployment of a DOCSIS system using the M-CMTS architecture.

A list of the documents in the Modular CMTS Interface Specifications family is provided below.

Designation	Title
SCTE 137-2 2010	M-CMTS Downstream External PHY Interface (this document)
SCTE 137-1 2010	DOCSIS Timing Interface
SCTE 137-4 2010	Edge Resource Manager Interface
SCTE 137-3 2010	M-CMTS Operations Support System Interface

1.2 Requirements and 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.

Throughout this 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 behavior is acceptable or even useful, but the full implications should be understood and the case carefully weighed before implementing any behavior described with this label.
"MAY"	This word 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.

2 REFERENCES

2.1 Normative References

In order to claim compliance with this specification, it is necessary to conform to the following standards and other works as indicated, in addition to the other requirements of this specification. Notwithstanding, intellectual property rights may be required to use or implement such normative references. At the time of publication, the editions indicated were valid. All references are subject to revision; users of this specification are therefore encouraged to investigate the possibility of applying the most recent edition of the standards and other references listed below.

[DRFI]	ANSI/SCTE 133 2010, Downstream RF Interface for Cable Modem Termination Systems
[DTI]	ANSI/SCTE 137-1 2010, Modular Head End Architecture Part 1: DOCSIS Timing Interface.
[ERMI]	ANSI/SCTE 137-4 2010, Modular Head End Architecture Part 4: Edge Resource Manager Interface.
[IANA-PORTS]	IANA, Port Numbers, June 2004.
[IEEE-802.1Q]	IEEE Std 802.1Q™-2003, Virtual Bridged Local Area Networks, May 2003.
[IEEE-802.3]	IEEE Std 802.3™-2002, Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications, March 2002.
[ISO 13818]	ISO/IEC 13818-1, INFORMATION TECHNOLOGY - GENERIC CODING OF MOVING PICTURES AND ASSOCIATED AUDIO: SYSTEMS Recommendation H.222.0, February 2000.
[ITU-T J.83]	ITU-T Recommendation J.83 (4/97), Digital multi-programme systems for television sound and data services for cable distribution.
[ITU-T Y.1541]	ITU-T Recommendation Y.1541 (05/2002), Internet protocol aspects – Quality of service and network performance.
[M-OSSI]	ANSI/SCTE 137-3 2010, Modular Head End Architecture Part 3: M-CMTS Operations Support System Interface.
[RFC 308]	IETF RFC 3308, Layer Two Tunneling Protocol (L2TP) Differentiated Services Extension, November 2002.
[RFC 791]	IETF RFC 791, Internet Protocol-DARPA, September 1981.
[RFC 1191]	IETF RFC 1191, MTU Path Discovery, November 1990.
[RFC 1981]	IETF RFC 1981, Path MTU Discovery for IP version 6, August 1996.
[RFC 2597]	IETF RFC 2597, Assured Forwarding PHB Group, June 1999.
[RFC 2983]	IETF RFC 2983, Differentiated Services and Tunnels, October 2000.
[RFC 3246]	IETF RFC 3246, An Expedited Forwarding PHB (Per-Hop Behavior), March 2002.
[RFC 3260]	IETF RFC 3260, New Terminology and Clarifications for Diffserv, April 2002.
[RFC 3931]	IETF RFC 3931, Layer Two Tunneling Protocol - Version 3 (L2TPv3), March 2005.
[RFC 768]	IETF RFC 768, User Datagram Protocol, August 1980.
[RFI2.0]	ANSI/SCTE 79-1 2009, DOCSIS 2.0 Part 1: Radio Frequency Interface.

2.2 Informative References

This document uses the following informative references:

[IANA-L2TP]	IANA, Layer Two Tunneling Protocol (L2TP) Parameters.
[ISO8802-2]	ISO/IEC 8802-2: 1994 (IEEE Std 802.2: 1994) - Information technology – Telecommunications and information exchange between systems - Local and metropolitan area networks - Specific requirements - Part 2: Logical link control.
[RFC 3140]	IETF RFC 3140, Per Hop Behavior Identification Codes, June 2001.
[VCCV]	IETF draft-ietf-pwe3-vccv-10.txt, T Nadeau, Pseudo Wire Virtual Circuit Connectivity Verification (VCCV), June 2006.

2.3 Reference Acquisition

- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027; Phone +1-303-661-9100; Fax +1-303-661-9199; <http://www.cablelabs.com>
- The Institute of Electrical and Electronics Engineers, Inc, Internet: <http://standards.ieee.org>
- Internet Assigned Numbers Authority, IANA, Internet: <http://www.iana.org>
- European Telecommunications Standards Institute, ETSI, <http://www.etsi.org>
- Internet Engineering Task Force (IETF) Secretariat, 46000 Center Oak Plaza, Sterling, VA 20166, Phone +1-571-434-3500, Fax +1-571-434-3535, <http://www.ietf.org>
- International Telecommunication Union (ITU), Phone +41 22 730 5111 (ITU Switchboard), Fax +41 22 733 7256, <http://www.itu.int/>
- International Organization for Standardization (ISO), Tel.: +41 22 749 02 22, Fax: +41 22 749 01 55, www.standardsinfo.net

3 TERMS AND DEFINITIONS

This specification uses the following terms:

Bonded Channels	A logical channel comprising multiple individual channels.
Cable Modem (CM)	A modulator-demodulator at subscriber locations intended for use in conveying data communications on a cable television system.
Converged Interconnect Network	The network (generally gigabit Ethernet) that connects an M-CMTS Core to an EQAM.
Customer Premises Equipment (CPE)	Equipment at the end user's premises; may be provided by the service provider.
Data Rate	Throughput, data transmitted in units of time usually in bits per second (bps).
Decibels (dB)	Ratio of two power levels expressed mathematically as $\text{dB} = 10\log_{10}(\text{P}_{\text{OUT}}/\text{P}_{\text{IN}})$.
Decibel-Millivolt (dBmV)	Unit of RF power expressed in decibels relative to 1 millivolt, where $\text{dBmV} = 20\log_{10}(\text{value in mV}/1 \text{ mV})$.
Downstream (DS)	1. Transmissions from CMTS to CM. This includes transmission from the M-CMTS Core to the EQAM, as well as the RF transmissions from the EQAM to the CM. 2. RF spectrum used to transmit signals from a cable operator's headend or hub site to subscriber locations.
Edge QAM modulator (EQAM)	A head end or hub device that receives packets of digital video or data. It re-packetizes the video or data into an MPEG transport stream and digitally modulates the digital transport stream onto a downstream RF carrier using quadrature amplitude modulation (QAM).
Flow	A stream of packets in DEPI used to transport data of a certain priority from the M-CMTS Core to a particular QAM channel of the EQAM. In PSP operation, there can exist several flows per QAM channel.
Gbps	Gigabits per second
Gigahertz (GHz)	A unit of frequency; 1,000,000,000 or 10^9 Hz.
GigE (GE)	Gigabit Ethernet (1 Gbps)
Hertz (Hz)	A unit of frequency; formerly cycles per second.
Hybrid Fiber/Coax (HFC) System	A broadband bidirectional shared-media transmission system using optical fiber trunks between the head-end and the fiber nodes, and coaxial cable distribution from the fiber nodes to the customer locations.
Institute of Electrical and Electronic Engineers (IEEE)	A voluntary organization which, among other things, sponsors standards committees and is accredited by the American National Standards Institute (ANSI).
Internet Engineering Task Force (IETF)	A body responsible for, among other things, developing standards used in the Internet.
Internet Protocol (IP)	An Internet network-layer protocol
kilohertz (kHz)	Unit of frequency; 1,000 or 10^3 Hz; formerly kilocycles per second
L2SS	Layer 2 Specific Sublayer. DEPI is a L2SS of L2TPv3.
L2TP Access Concentrator (LAC)	If an L2TP Control Connection Endpoint (LCCE) is being used to cross-connect an L2TP session directly to a data link, we refer to it as an L2TP Access Concentrator (LAC). An LCCE may act as both an L2TP Network Server (LNS) for some sessions and an LAC for others, so these terms must only be used within the context of a given set of sessions unless the LCCE is, in fact, single purpose for a given topology.

L2TP Attribute Value Pair (AVP)	The L2TP variable-length concatenation of a unique Attribute (represented by an integer), a length field, and a Value containing the actual value identified by the attribute.
L2TP Control Connection	An L2TP control connection is a reliable control channel that is used to establish, maintain, and release individual L2TP sessions, as well as the control connection itself.
L2TP Control Connection Endpoint (LCCE)	An L2TP node that exists at either end of an L2TP control connection. May also be referred to as an LAC or LNS, depending on whether tunneled frames are processed at the data link (LAC) or network layer (LNS).
L2TP Control Connection ID	The Control Connection ID field contains the identifier for the control connection, a 32-bit value. The Assigned Control Connection ID AVP, Attribute Type 61, contains the ID being assigned to this control connection by the sender. The Control Connection ID specified in the AVP must be included in the Control Connection ID field of all control packets sent to the peer for the lifetime of the control connection. Because a Control Connection ID value of 0 is used in this special manner, the zero value must not be sent as an Assigned Control Connection ID value.
L2TP Control Message	An L2TP message used by the control connection.
L2TP Data Message	L2TP message used by the data channel
L2TP Endpoint	A node that acts as one side of an L2TP tunnel
L2TP Network Server (LNS)	If a given L2TP session is terminated at the L2TP node and the encapsulated network layer (L3) packet processed on a virtual interface, we refer to this L2TP node as an L2TP Network Server (LNS). A given LCCE may act as both an LNS for some sessions and an LAC for others, so these terms must only be used within the context of a given set of sessions unless the LCCE is in fact single purpose for a given topology.
L2TP Pseudowire (PW)	An emulated circuit as it traverses a packet-switched network. There is one Pseudowire per L2TP Session.
L2TP Pseudowire Type	The payload type being carried within an L2TP session. Examples include PPP, Ethernet, and Frame Relay.
L2TP Session	An L2TP session is the entity that is created between two LCCEs in order to exchange parameters for and maintain an emulated L2 connection. Multiple sessions may be associated with a single Control Connection.
L2TP Session ID	A 32-bit field containing a non-zero identifier for a session. L2TP sessions are named by identifiers that have local significance only. That is, the same logical session will be given different Session IDs by each end of the control connection for the life of the session. When the L2TP control connection is used for session establishment, session IDs are selected and exchanged as Local Session ID AVPs during the creation of a session. The Session ID alone provides the necessary context for all further packet processing, including the presence, size, and value of the Cookie, the type of L2-Specific Sublayer, and the type of payload being tunneled.
MAC Domain	A grouping of layer 2 devices that can communicate with each other without using bridging or routing. In DOCSIS is the group of CMs that are using upstream and downstream channels linked together through a MAC forwarding entity.
Maximum Transmission Unit (MTU)	The layer 3 payload of a layer 2 frame.
Mbps	Megabits per second
Media Access Control (MAC)	Used to refer to the layer 2 element of the system which would include DOCSIS framing and signaling.

Megahertz (MHz)	A unit of frequency; 1,000,000 or 10^6 Hz; formerly megacycles per second
Microsecond (μs)	10^{-6} second
Millisecond (ms)	10^{-3} second
Modulation Error Ratio (MER)	The ratio of the average symbol power to average error power
Multiple System Operator (MSO)	A corporate entity that owns and/or operates more than one cable system.
Nanosecond (ns)	10^{-9} second
Physical Media Dependent (PMD) Sublayer	A sublayer of the Physical layer which is concerned with transmitting bits or groups of bits over particular types of transmission link between open systems and which entails electrical, mechanical, and handshaking procedures.
QAM channel (QAM ch)	Analog RF channel that uses quadrature amplitude modulation (QAM) to convey information.
Quadrature Amplitude Modulation (QAM)	A modulation technique in which an analog signal's amplitude and phase vary to convey information, such as digital data.
Radio Frequency (RF)	In cable television systems, this refers to electromagnetic signals in the range 5 to 1000 MHz.
Radio Frequency Interface (RFI)	Term encompassing the downstream and the upstream radio frequency interfaces.
Request For Comments (RFC)	A technical policy document of the IETF; these documents can be accessed on the World Wide Web at http://www.rfc-editor.org/ .
Session	An L2TP data plane connection from the M-CMTS Core to the QAM channel. There must be one session per QAM Channel. There is one DEPI pseudowire type per session. There may be one MPT flow or one or more PSP flows per session. Multiple sessions may be bound to a single control connection.
StopCCN	L2TPv3 Stop-Control-Connection-Notification message
Upconverter	A device used to change the frequency range of an analog signal, usually converting from a local oscillator frequency to an RF transmission frequency.
Upstream (US)	1. Transmissions from CM to CMTS. This includes transmission from the EQAM to M-CMTS Core as well as the RF transmissions from the CM to the EQAM. 2. RF spectrum used to transmit signals from a subscriber location to a cable operator's headend or hub site.
Upstream Channel Descriptor (UCD)	The MAC Management Message used to communicate the characteristics of the upstream physical layer to the cable modems.
Video on Demand (VoD) System	System that enables individuals to select and watch video content over a network through an interactive television system.

4 ABBREVIATIONS AND ACRONYMS

This specification uses the following abbreviations:

ACK	L2TPv3 Explicit Acknowledgement message
AVP	L2TPv3 Attribute Value Pair
CDN	L2TPv3 Call-Disconnect-Notify message
CIN	Converged Interconnect Network
CLI	Command Line Interface
CM	Cable Modem
CMCI	Cable Modem CPE Interface
CMTS	Cable Modem Termination System
CPE	Customer Premises Equipment
CRC	Cyclic Redundancy Check
CRC16	CRC of length 16
CSMA	Carrier Sense Multiple Access
dB	Decibels
dBmV	Decibel-Millivolt
DEPI	Downstream External-PHY Interface
DOCSIS®	Data-Over-Cable Service Interface Specifications
DOCSIS-MPT (D-MPT)	DOCSIS MPT Mode
DRFI	Downstream Radio Frequency Interface
DS	Downstream
DSCP	Differentiated Services Code Point
DTI	DOCSIS Timing Interface
DTS	DOCSIS Time Stamp, 32-bit
EQAM	Edge QAM
ERM	Edge Resource Manager
ERMI	Edge Resource Manager Interface
ETSI	European Telecommunications Standards Institute
FQDN	Fully Qualified Domain Name
Gbps	Gigabits per second
GHz	Gigahertz
GE	Gigabit Ethernet (Gig E)
Hz	Hertz
HELLO	L2TPv3 Hello message
HFC	Hybrid Fiber/Coax
ICCN	L2TPv3 Incoming-Call-Connected message
ICRP	L2TPv3 Incoming-Call-Reply message
ICRQ	L2TPv3 Incoming-Call-Request message
IEEE	Institute of Electrical and Electronic Engineers
IETF	Internet Engineering Task Force

IP	Internet Protocol
IPv4	Internet Protocol version 4
ISO	International Standards Organization
ITU	International Telecommunications Union
ITU-T	Telecommunication Standardization Sector of the International Telecommunication Union
kbps	Kilobits per second
kHz	Kilohertz
L2TP	Layer 2 Transport Protocol
L2TPv3	Layer 2 Transport Protocol – version 3
L3	Layer 3
LAC	L2TP Access Concentrator
LCCE	L2TP Control Connection Endpoint
LNS	L2TP Network Server
LSB	Least Significant Bit
MAC	Media Access Control
Mbps	Megabits per second
M-CMTS™	Modular Cable Modem Termination System
MER	Modulation Error Ratio
MHz	Megahertz
MIB	Management Information Base
M/N	Relationship of integer numbers M,N that represents the ratio of the downstream symbol clock rate to the DOCSIS master clock rate
MPEG	Moving Picture Experts Group
MPEG-TS	Moving Picture Experts Group Transport Stream
MPT	MPEG-TS mode of DEPI
MPTS	Multi Program Transport Stream
ms	Millisecond
MSO	Multiple System Operator
MSB	Most Significant Bit
MTU	Maximum Transmission Unit
ns	Nanosecond
OSSI	Operations System Support Interface
PCR	Program Clock Reference. A time stamp in the Video Transport Stream from which decoder timing is derived.
PHY	Physical Layer
PID	Packet Identifier; PID (system): A unique integer value used to identify elementary streams of a program in a single or multi-program Transport Stream as described in 2.4.3 of ITU-T Rec. H.222.0 [ISO 13818]
PMD	Physical Media Dependent Sublayer
PPP	Point-to-Point Protocol
PSI	Program Specific Information
PSP	Packet-Streaming-Protocol

PW	Pseudowire
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RFI	Radio Frequency Interface
RFC	Request For Comments
SCCRN	L2TPv3 Start-Control-Connection-Connected message
SCCRP	L2TPv3 Start-Control-Connection-Reply message
SCCRQ	L2TPv3 Start-Control-Connection-Request message
S-CDMA	Synchronous Code Division Multiple Access
SLI	L2TPv3 Set Link Info message
SPTS	Single Program Transport Stream
StopCCN	L2TPv3 Stop-Control-Connection-Notification message
TSID	MPEG2 Transport Stream ID
UCD	Upstream Channel Descriptor
UDP	User Datagram Protocol
US	Upstream
VOD	Video On Demand

5 TECHNICAL OVERVIEW

This section is informative.

5.1 System Architecture




Figure 5-1 - Modular CMTS Reference Architecture

In the M-CMTS architecture, a device referred to as the M-CMTS Core contains the DOCSIS MAC. This includes all signaling functions, downstream bandwidth scheduling, and DOCSIS framing. The EQAM box contains mainly PHY related circuitry, such as QAM modulators, and tunneling logic, to connect to the M-CMTS Core.

5.1.1 Reference Architecture

The reference architecture for a Modular CMTS system is shown in Figure 5-1. This architecture contains several pieces of equipment along with interfaces between those pieces of equipment. This section will briefly introduce each device and interface.

The Edge QAM device, or EQAM for short, has its origins in the VOD environment. It is a chassis that typically has one or more gigabit Ethernet coming in and multiple QAM modulators and RF upconverters on the output. This EQAM is being adapted for use in a Modular CMTS environment. The individual outputs of these devices are often referenced to as QAM Channel rather than the full "QAM Modulator and RF Upconverter".

The M-CMTS Core contains everything a traditional CMTS does, except for functions performed in the EQAM. The M-CMTS Core contains the downstream MAC and all the initialization and operational DOCSIS related software.

This diagram currently shows the Upstream Receivers of DOCSIS upstreams located internally to the M-CMTS Core. However, there is nothing preventing an implementation of a Modular CMTS from using external upstream receivers. In the future, upstream receivers may be external to the M-CMTS Core.

The DOCSIS Timing Interface (DTI) Server provides a common frequency of 10.24 MHz and a DOCSIS timestamp to other M-CMTS elements.

DEPI, the Downstream External PHY Interface, is the interface between the M-CMTS Core and the EQAM. More specifically, it is an IP Tunnel between the MAC and PHY in a Modular CMTS system which contains both a data path for DOCSIS frames, and a control path for setting up, maintaining, and tearing down sessions.

DRFI, or Downstream Radio Frequency Interface, is intended to capture all the current and future RF requirements for the downstream direction for both integrated DOCSIS CMTS systems, Modular DOCSIS CMTS systems, and VOD EQAM systems.

DTI, or DOCSIS Timing Interface, is a point-to-point interface from the DTI Server to other M-CMTS elements. The DTI Specification [DTI] defines DTI Server and DTI Client behaviors and protocols. The DTI Server is the Timing Signal Generator while each M-CMTS Core and EQAM has a DTI Client. The DTI Server distributes a 10.24 MHz frequency and a DOCSIS timestamp over unshielded twisted pair (UTP). The DTI protocol automatically compensates for cable length and ensures that all M-CMTS elements have the same sense of time and frequency.

ERMI, or Edge Resource Manager Interface [ERMI], involves three interfaces: a registration interface between an EQAM and ERM (Edge Resource Manager), a control interface between an EQAM and an ERM, and a control interface between an M-CMTS Core and an ERM. The first interface is used to register and unregister EQAM resources (i.e., QAM channels) with an ERM. The second interface is used by an ERM to request QAM channel resources from an EQAM, and by an EQAM to deliver resources to an ERM. The third interface is used by the M-CMTS Core to request specific QAM channel resources from the ERM, and by the ERM to respond to such requests with the location of QAM channel resources.

MOSSI, or Modular CMTS Operations Support System Interface [M-OSSI], provides the management interface to each system component. This interface is an extension of the OSSI defined in the DOCSIS specifications for monitoring a management of CMTS functions. This interface could be used, in place of an ERM and the ERMI, to statically configure and associate QAM channel resources with M-CMTS Cores. This interface allows for the modification of a QAM channel's physical layer parameter by either the M-CMTS Core or the EQAM, and provides a mechanism by which the operator can "lock" certain parameters at the EQAM so that they can only be modified there. This document defines the mechanism to communicate these parameter settings to the other side.

NSI, or the Network Side Interface, is unchanged, and is the physical interface the CMTS uses to connect to the backbone network. Today, this is typically 100 Mbps or 1 Gbps Ethernet.

CMCI, or Cable Modem to Customer Premise Equipment Interface, is also unchanged, and is typically 10/100 Mbps Ethernet or USB.

5.1.2 DEPI Operation

DEPI is an IP Tunnel that exists between the DOCSIS MAC in the M-CMTS Core and the DOCSIS PHY that exists in the EQAM. DEPI's job is to take either formatted DOCSIS frames or MPEG packets and transport them through a layer 2 or layer 3 network and deliver them to the EQAM for transmission.

The base protocol that is used for the DEPI is the Layer 2 Tunneling Protocol Version 3, or L2TPv3 for short [RFC 3931]. L2TPv3 is an IETF protocol that is a generic protocol for creating a pseudowire. A pseudowire is a mechanism to transparently transport a layer 2 protocol over a layer 3 network. Examples of protocols supported by L2TPv3 include ATM, HDLC, Ethernet, Frame Relay, PPP, etc.

Section 8.1 "L2TPv3 Transport Packet Format" shows the format of an L2TPv3 data packet. Each data packet contains a 32-bit session ID which is associated with a single QAM Channel. The UDP header is optional in the L2TPv3 protocol.

L2TPv3 then permits a sub-header to exist whose definition is specific to the payload being carried. The control channel allows for signaling messages to be sent between the M-CMTS Core and EQAM. Typical control messages will set up a "control connection" between the M-CMTS Core and EQAM, and then set up multiple data sessions (one for each downstream QAM channel). Each session can be marked with different DiffServ Code Points (DSCPs) and support different encapsulation protocols.

There are two basic tunneling techniques defined by DEPI. The first technique, known as D-MPT mode, transports multiple 188-byte MPEG-TS packets by placing them into the L2TPv3 payload with a unique sub-header which contains a sequence number so packet drops can be detected. The encapsulation of DOCSIS Frames into MPEG-TS packets is performed in the M-CMTS Core. The second technique, known as the Packet Streaming Protocol (PSP), transports DOCSIS Frames in the L2TPv3 payload. The DOCSIS Frames are then encapsulated in MPEG-TS packets within the EQAM. PSP mode allows DOCSIS frames to be both concatenated, to increase network performance, and fragmented, in case the tunneled packets exceed the network MTU size.

One of the technical considerations of the Modular CMTS architecture is its impact on the round trip request-grant delay time. The request-grant delay time is the time from when a CM requests bandwidth using an uncontended bandwidth request (REQ) to when it receives a MAP message with the granted transmit opportunity in it.

To prevent the MAP from being slowed down by other traffic in the CIN, the DOCSIS traffic (or a subset containing the MAP messages) may be sent in an independent L2TPv3 flow that has a unique DSCP. The value of the marked DSCP value should be consistent with a configured "per hop behavior (PHB)" that will provide MAP messages with the highest priority and lowest latency across the CIN to the EQAM.

5.1.3 EQAM Operation




Figure 5-2 - EQAM Block Diagram

Figure 5-2 shows a high level block diagram of an EQAM that is capable of handling either Video MPEG traffic or DOCSIS traffic. The expression "D-MPT" is an acronym for DOCSIS MPEG Transport.

The first interface that is shown is the VOD transport. VOD SPTS or MPTS streams are received with a format of MPEG packets over UDP/IP. The video processing functions generally include de-jittering, re-multiplexing PID remapping, MPEG-2 PSI insertion, and PCR timestamp correction. These functions are not defined in this specification.

The next set of interfaces are the DEPI interfaces.

The first interface defined is D-MPT. This is a mode where the EQAM must search the incoming D-MPT frames for DOCSIS SYNC messages and correct the timestamp value in these messages based upon the EQAM's internal DOCSIS timestamp, which has been derived from DTI. The resulting D-MPT frames are then copied to the QAM channel without further interpretation or modification. This mode is intended for DOCSIS frames where the MAP is embedded into the stream and network latency or jitter is not a concern.

Because D-MPT mode encapsulates all DOCSIS traffic into a single DEPI flow, it does not allow for QoS differentiation among various types of traffic either across the CIN or within the EQAM. For example, acceleration of MAPs relative to other DOCSIS data is not possible when D-MPT mode is used. This may still give acceptable performance when the delay and jitter introduced onto DEPI packets by the CIN and EQAM are determined by the operator to be sufficiently low.

Certain network conditions and/or architectures may reduce network delay and/or jitter and make it more likely that DOCSIS MPT mode will give acceptable system performance. Some examples are:

- networks with a very small number of hops (e.g., 1 or 2 hops)
- networks which are largely shared with video traffic from VOD servers to EQAMs, on which most of the traffic is video and it is acceptable to prioritize all DEPI traffic on the network above all VOD traffic
- networks which are lightly loaded so that delays due to congestion are highly unlikely

Such conditions may exist on today's CINs. They may be less likely to occur as deployments become denser and total IP traffic increases. Evaluation of network conditions and decisions about acceptable performance levels with DOCSIS MPT mode are in the realm of operator discretion.

The next interface is DOCSIS PSP. This interface transports DOCSIS Data and MAPs in separate flows which are streamed together in a uniform byte stream by the M-CMTS Core. The PSP re-assembly engine removes this overhead and recovers the DOCSIS Frames. The PSP scheduler then allows MAPs to be placed in order, ahead of data and SYNC messages. In PSP mode, the EQAM generates all SYNC messages as detailed in Section 6.1.1. The output is then delivered to a transmission convergence layer which converts the results to a DOCSIS MPEG stream.

The last interface is the DTI interface which provides a common frequency and DOCSIS timestamp. The reference is used to synchronize the downstream symbol rate and DOCSIS timestamp for use with DOCSIS cable modems. The timestamp is used for the DOCSIS SYNC correction.

The output from the EQAM consists of a stream of MPEG frames which carry video and/or DOCSIS data, and are modulated onto an RF carrier in accordance with the DRFI specification [DRFI].

5.2 Bonding Services Model




Figure 5-3 - Bonding Services Model

Downstream Channel Bonding refers to the forwarding of a stream of DOCSIS frames across multiple QAM carriers.

In the Modular CMTS architecture, the Downstream Channel Bonding is implemented in the M-CMTS Core. At the M-CMTS Core, packets from the IP backbone are placed in a DOCSIS frame. That DOCSIS frame is then sent to one of several QAM channels in a Bonding Group. The frame may be transported using D-MPT or PSP.

In this system, the EQAM is unaware that it is doing bonding. It is also unaware of any of the details of the Bonding protocol.

5.3 Multiple Services Model




Figure 5-4 - Multi-Service Mode

The Modular CMTS (M-CMTS) architecture reflects an effort to merge the video-on-demand (VOD) and high-speed-data (HSD) applications. This effort is made in anticipation of achieving higher efficiencies than are possible with two separate transmission networks feeding into the cable plant. In particular, the M-CMTS effort has repartitioned the traditional CMTS architecture so that the transmission technology which is common to both the VOD environment and the HSD environment will be able to share the same EQAM devices.

In a video delivery system, the EQAM is used to deliver video streams in MPEG-TS format to set top boxes at the subscribers' premises. This functionality will continue to be used in the future and operates independently of other processing which is described here. The M-CMTS architecture adds new interfaces which are unique to HSD service. These interfaces support both traditional DOCSIS and multi-channel (bonded) DOCSIS payloads as received from the M-CMTS Core device.

The M-CMTS Core provides the gateway functionality between the IP-based core network and the CIN. As such, the M-CMTS Core provides support for a multitude of services – including, but not limited to video over IP, voice-over-IP (VoIP), email, gaming, video telephony, etc.

6 DEPI ARCHITECTURE

This section is normative.

6.1 DEPI Data Path




Figure 6-1 - Downstream EQAM Block Diagram

A simplified logical block diagram of the internal data path of the EQAM is shown in Figure 6-1.

It is recognized, but not specified by this document, that the EQAM may receive non-DOCSIS MPEG elementary streams which have been encapsulated in MPEG packets and placed in a UDP datagram. This specification does not define requirements for transport of this type of traditional MPEG video. However, it is recognized that M-CMTS EQAM implementations may (and likely will) be capable of transporting traditional MPEG video (either interleaved with DOCSIS traffic on a single QAM channel, or on separate QAM channels within the same EQAM chassis).

The M-CMTS Core MUST support PSP mode, D-MPT mode, or both. The EQAM MUST support PSP mode, D-MPT mode, or both.

Within a session, the M-CMTS Core MUST support either one priority level of D-MPT or at least two priority levels of PSP, where each priority level would have a different DSCP. The M-CMTS Core MUST provide a mechanism to map DOCSIS traffic to the multiple priority levels of PSP. The M-CMTS Core MUST NOT attempt to establish a session which includes both PSP and D-MPT flows.

Within a session, the EQAM MUST support either one priority level of D-MPT or at least two priority levels of PSP, where each priority level would have a different DSCP. The EQAM is not intended to simultaneously support D-MPT and PSP within a single session, and MUST reject any attempt to establish such a session. Each priority level for each DEPI type, maps to one or more DEPI flows. The EQAM MUST support the establishment of a single DEPI flow per priority level. The EQAM MAY support the establishment of more than one DEPI flow per priority

level. The mapping of individual flows to priority levels (QoS Queues) is vendor-specific (see Figure 7–3). The mapping of flows will be done through the local EQAM Command Line Interface (CLI) configuration.

For both D-MPT and PSP modes, the EQAM **MUST** insert null MPEG packets when it has no data to send. The EQAM **SHOULD NOT** insert a null MPEG packet if it has data to send. Note that MPEG null insertion should take place prior to the correction of the DOCSIS SYNC message (refer to Section 6.1.1).

6.1.1 DOCSIS D-MPT Data Path

The DEPI DOCSIS D-MPT flows contain DOCSIS frames using the format as described in Section 8.2. All DOCSIS frames, including packet based frames and MAC management based frames, are included within the one D-MPT flow. The EQAM searches the D-MPT payload for any DOCSIS SYNC messages and performs SYNC corrections as described in Section 6.1.3.2. It then forwards the D-MPT packet to the RF interface.

The intent of D-MPT mode is that MPEG packets can be received by the EQAM and forwarded directly to the RF interface without having to terminate and regenerate the MPEG framing. The only manipulation of the D-MPT payload is the SYNC correction.

6.1.2 PSP Data Path

The Packet Stream Protocol (PSP) is a layer-3 convergence layer protocol, which allows packets to be consecutively streamed together and fragmented at arbitrary boundaries. The intent of the PSP mode is to facilitate Quality of Service. This mode is to be used for transporting traditional DOCSIS data and signaling messages which use one or more DSCP values. For example, in order to reduce Request-Grant latency, MAP MAC management messages may be sent using a different DSCP on a different PSP flow than the rest of the DOCSIS channel. Refer to Section 6.2.1 for more information. The EQAM **MUST** support a minimum of two PSP receivers per QAM modulator. The intent of two receivers is to permit the implementation of a higher latency PSP flow and a lower latency PSP flow.

Each PSP flow is terminated, and the DOCSIS Frames within the flow are extracted. The DOCSIS frames are placed into corresponding output QoS queues. The output of the QoS queues go to a Packet Scheduler, which decides which queue is to be serviced based upon the PHB (negotiated between the M-CMTS Core and EQAM) of the PSP flow, which carried the DOCSIS frames. The Packet Scheduler is also responsible for inserting DOCSIS SYNC messages within the time interval specified by the DOCSIS SYNC Control AVP (see Figure 7–31). The Packet Scheduler **SHOULD** support a strict priority scheduler. The Packet Scheduler **MAY** support other queue scheduling disciplines.

The phrase "packet scheduler" is a general term that describes a method of applying priorities to different queues as packets are moved from different input queues to the output queue. An example of a typical Packet Scheduling algorithm would be Weighted Fair Queuing (WFQ) where some streams are given priority over other streams, but only up to some limit. This should not be confused with the more complex DOCSIS upstream scheduler, which deals with requests and grants.

The output of the Packet Scheduler goes to a Transmission Convergence engine which places the DOCSIS frames into MPEG packets according to the requirements in [DRFI]. This includes the insertion of stuffing bytes and the DOCSIS SYNC message as described by Section 6.1.3. The output of Transmission Convergence (TC) engine is sent to the RF Interface.

PSP mode primarily provides acceleration of MAPs through the network, in an attempt to reduce Request Grant latency. The PSP mode is most useful when all or most traffic has migrated to DOCSIS, and therefore, marking MAPs with higher QoS, when compared to other DOCSIS traffic that is needed in order to provide lower latency for MAPs, when traversing a fully subscribed network. Therefore, PSP has value in the long term and is included to address the case where most or all traffic is carried to the home via DOCSIS.

6.1.3 DOCSIS SYNC Message

6.1.3.1 SYNC Message Format

The DOCSIS Time Synchronization (SYNC) MAC message is transmitted by a Modular CMTS system at a periodic interval to establish MAC sublayer timing in cable modems. This message **MUST** use an FC field with FC_TYPE =

MAC Specific Header and FC_PARM = Timing MAC Header. This MUST be followed by a Packet PDU in the format shown in Figure 6–2.




Figure 6–2 - Format of a DOCSIS SYNC MAC Message

The fields MUST be as defined below:

FC, MAC PARM, LEN, HCS: Common MAC frame header with FC_PARM field to indicate a timing header – refer to [RFI2.0] for details.

Destination Address (DA): Set to the DOCSIS MAC multicast address of 01-E0-2F-00-00-01.

Source Address (SA): The MAC address of the M-CMTS Core. In PSP Mode, the EQAM learns the appropriate M-CMTS Core MAC address via explicit signaling during L2TPv3 session setup.

Msg Length: Length of the MAC message from DSAP to the end of the payload.

DSAP: The LLC null destination SAP (00) as defined by [ISO8802-2].

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

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

DOCSIS Version: Set to 1.

DOCSIS Type: Set to 1 to indicate a SYNC message.

CMTS Timestamp: The count state of an incrementing 32-bit binary counter clocked with the 10.24 MHz master clock derived from [DTI].

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 [DRFI].

6.1.3.2 Correction and Insertion

The EQAM MUST derive a local DOCSIS timestamp from the DTI Client, specified in [DTI]. In the D-MPT mode, the EQAM must be capable of correcting all embedded SYNC messages in the DOCSIS stream. For the DOCSIS PSP mode, the EQAM MUST be capable of inserting DOCSIS SYNC messages based upon its internal timestamp into the downstream MPEG-TS stream. In PSP Mode, the EQAM MUST insert the SYNC message starting at the 6th byte of the MPEG-TS frame (the fifth byte would be the MPEG pointer_field). When referencing a timebase to use for SYNC timestamp insertion or correction, the EQAM MUST utilize a timebase that is delayed by no less than

0 and no more than 100 clock cycles (approximately 10 μ s) of the master clock (10.24 MHz) compared to the time communicated via the DTI Client Test Port output.

This time difference between the time communicated via the DTI Client Test Port output and the time reference applied by an EQAM is expected to be essentially constant. All timing and jitter requirements that exist in this specification and [DRFI] still apply. The requirements here do not preclude an EQAM from taking longer than 100 clock cycles to process the DTI conveyed time, but in that case, the EQAM would need to internally adjust the applied timebase such that it is delayed between zero and 100 clock cycles from the DTI conveyed timebase.

When using the D-MPT mode, the M-CMTS Core MUST generate SYNC messages and include them in the D-MPT payload. The M-CMTS Core MUST insert the SYNC message starting at the 6th byte of the MPEG-TS frame (the fifth byte would be the MPEG pointer_field). This is intended to simplify the EQAM implementation by allowing the EQAM to check only the payload_unit_start_indicator bit and the fifth and sixth byte (which together will contain 0x00C0) of the MPEG-TS packet to locate a DOCSIS SYNC message. Note that the CMTS Timestamp in the SYNC messages generated by the M-CMTS core is not required to accurately reflect the current timestamp received via a DTI Client in the M-CMTS Core. For example, it is permissible for the M-CMTS Core to use a value of zero for the CMTS Timestamp in all SYNC messages. When using the DOCSIS PSP mode, the M-CMTS Core MUST NOT generate SYNC messages as part of the PSP payload.

6.1.4 Latency and Skew Requirements

6.1.4.1 Latency

For PSP DEPI sessions, latency is defined as the absolute difference in time from when the last bit of a DEPI packet containing the last bit of a single DOCSIS MAC frame enters the EQAM DEPI port to the time that the first bit of the DOCSIS MAC frame exits the EQAM RFI port. For D-MPT DEPI sessions, latency is defined as the absolute difference in time from when the last bit of a DEPI packet enters the EQAM DEPI port to the time that the first bit of the first MPEG packet contained within said DEPI packet exits the EQAM RFI port. At the EQAM input, the last bit of the arriving DEPI packet is used because the EQAM's layer 2 interface (e.g., GigE Ethernet port) must receive the entire packet before the EQAM can begin processing. At the EQAM output, the first bit of the first MPEG packet (in D-MPT mode) or DOCSIS frame (in PSP mode) inside the DEPI packet is used to guarantee that the data complies with the definition of "isolated packets" (see next paragraph). If this were not done, the measurement could be corrupted by delays incurred due to queuing of data behind other packets destined for the same RF interface. The EQAM SHOULD allow enough buffer memory per QAM channel to buffer at least 20 ms worth of traffic across all L2TPv3 sessions destined for that QAM channel.

In PSP DEPI sessions, the multiple flows supported in the EQAM provide for prioritized access to the modulator. In the absence of higher priority traffic, and regardless of load of lower priority traffic, the EQAM MUST forward isolated packets in each DEPI flow with a latency of less than 500 μ s plus the delay of the interleaver. Isolated packets are spaced such that, at the nominal downstream data rate, the EQAM would complete transmission of the current packet before the arrival of the next packet.

6.1.4.2 Skew

Skew is defined as the difference between the maximum latency and the minimum latency through the EQAM, as measured from two reference bits on the network interface to the same bits on two separate RF outputs. Skew is to be measured with the PHY parameters set equal on the QAM Channels under measurement.

The skew between any two bonded QAM Channels within an EQAM MUST be less than 500 μ sec. The skew requirements for the EQAM are implicitly met when the EQAM meets the latency requirements set out above in Section 6.1.4.1. This requirement is intended for the transmission of skew sensitive traffic such as bonded traffic.

6.2 Networking Considerations

6.2.1 Per Hop Behavior Usage

The IETF has defined a number of Per Hop Behaviors (PHBs) to be used for offering network-based QoS. DEPI supports use of the 6-bit Expedited Forwarding (EF) PHB as described in [RFC 3246], Assured Forwarding (AF) PHBs as described in [RFC 2597], and best effort forwarding as described in [RFC 2597]. DEPI negotiates six-bit

Per-Hop Behavior Identifiers (PHBIDs) between the M-CMTS Core and the EQAM. Note: the term PHBID, as defined in this specification, while consistent with the purpose of PHBIDs defined in [RFC 3140], does not conform to the format of PHBIDs defined in [RFC 3140]. The term "PHBID" in this specification is equivalent to the term "recommended DSCP value" in [RFC 3140]. The M-CMTS Core MUST support the Expedited Forwarding PHBID for PSP mode. The M-CMTS Core MUST support the best effort PHBID for both DEPI modes. The EQAM MUST support the EF PHBID for PSP mode. The EQAM MUST support the best effort PHBID for both DEPI modes. Table 6–1 lists the PHBs explicitly supported by the DEPI specification. Note: This specification does not preclude support for other PHBs not defined in Table 6–1.

Table 6–1 - PHBs and recommended DSCP values

PHB	PHB ID(s) and recommended DSCP value(s)
EF	46
AF (multiple levels)	10, 12, 14, 18, 20, 22, 26, 28, 30, 34, 36, 38
Best effort	0

The DEPI interface supports multiple traffic types including DOCSIS MAC and DOCSIS data traffic. Within both traffic types, there may be different levels of priority. For PSP operation, the M-CMTS Core SHOULD provide a mechanism to map traffic of different priorities to DEPI flows with different PHB values. The M-CMTS Core SHOULD NOT use the same PHB across multiple DEPI flows within a session.

The CIN should provide the appropriate Per Hop Behavior for the differentiated traffic types. The level of granularity provided for differentiated traffic is determined by the network operator, but at a minimum, it would be expected that DOCSIS MAP messages and VoIP data traffic are prioritized over best effort data traffic.

The EQAM uses the PHB signaled in the establishment of the DEPI flow when scheduling multiple DEPI PSP flows onto one QAM Channel as described in Section 6.1.2.

6.2.2 DiffServ Code Point Usage

Differentiated Services (DiffServ) is a mechanism for tagging IP packets with defined PHBs. DiffServ uses the first 6 bits of the ToS Byte in DiffServ Field of the IP Header. The PHB values encoded into the DiffServ Field are referred to as DiffServ Code Points (DSCPs). The DSCP of the L2TPv3 packet SHOULD be assigned at the egress of the M-CMTS Core to provide Quality of Service for DEPI traffic within the CIN. The M-CMTS Core SHOULD use the DSCP values corresponding to the PHBs identified in Table 6–1. The DSCP MAY be used at the ingress of the EQAM.

The M-CMTS Core MUST tag all packets within a DEPI flow of an L2TPv3 session with the same DSCP.

DOCSIS frames encapsulated in L2TPv3 packets may contain IP packets which also have a DSCP assigned. The EQAM is not required to schedule packets based upon the original DSCP contained within the DOCSIS frame.

6.2.3 Packet Sequencing

For a stream of packets transmitted on a DEPI flow, the packet sequence number will be incremented by one for each packet sent, as described in Sections 8.2 and 8.3. If the EQAM detects a discontinuity in the packet sequence numbers indicating that one or more packets were dropped or delayed, an error is logged and the EQAM SHOULD transfer the current packet to the QAM Channel without waiting for the missing packets. If the EQAM detects a discontinuity in the packet sequence numbers indicating that one or more packets have arrived late, then those packets SHOULD be discarded. The EQAM MUST NOT forward packets which were skipped due to a discontinuity in the sequence numbers. Storing and re-ordering of packets so that they can be delivered to the QAM Channel in the correct sequence is not prohibited by these requirements and the EQAM MAY perform such re-ordering as long as the latency requirements of Section 6.1.4 are met.

6.2.4 Network MTU

The network between the M-CMTS Core and the EQAM will have a certain Maximum Transmission Unit. If a maximum size DOCSIS frame were to be tunneled from the M-CMTS Core to the EQAM without fragmentation,

the size of the resulting packet may be greater than the CIN can handle. Both the D-MPT and PSP modes avoid this issue by offering streaming and fragmentation. As such, IP fragmentation is not required. IP fragmentation is also undesirable because the EQAM may forward packets based upon the destination UDP port, and the UDP port is only available in the first IP fragment.

Determining the MTU to use for the L2TPv3 tunnel between the M-CMTS Core and the EQAM is a two step process. The first step is done as part of L2TPv3 session establishment (see Section 7) using the DEPI MTU AVPs. When the M-CMTS Core starts the session establishment with an ICRQ message it MUST supply the DEPI Local MTU AVP with a payload size that is the lesser of its receive capabilities and the receive capabilities defined by its lower layer. The receive capabilities of the M-CMTS Core are defined by its internal constraints, and any configured maximums. The receive capabilities defined by its lower layer is a calculation based on referencing the payload size constraints of the interface below which this tunnel is being created, as defined in Annex A.1. The M-CMTS Core MUST support an MTU size of at least 1500 bytes. The EQAM MUST send L2TPv3 frames with a payload size less than or equal to this maximum. If the EQAM cannot meet this criterion then it MUST fail session creation by generating a CDN message. The EQAM needs to consider the same criteria in calculating its MTU. The EQAM MUST support an MTU size of at least 1500 bytes, as calculated in Annex A.1. The EQAM MUST insert the DEPI Remote MTU AVP in the ICRP message with its MTU size. The M-CMTS Core MUST send L2TPv3 frames with a payload size less than or equal to this maximum. If the M-CMTS Core cannot meet this criterion then it MUST fail session creation by generating a CDN message.

The second step is to determine the MTU of the path between the M-CMTS Core and the EQAM. The M-CMTS Core MUST support a mechanism to prevent sending packets larger than the network MTU. This SHOULD be done using Path MTU Discovery, as described in [RFC 1191]. Annex A.3. gives a brief overview of the Path MTU discovery protocol. Alternatively, this MAY be done via a static configuration option. Both the M-CMTS Core and the EQAM MUST have a way to statically configure an MTU for each L2TPv3 session. To avoid IP fragmentation, the M-CMTS Core and the EQAM MUST set the Don't Fragment bit (DF) in the IPv4 header for all transmissions into the L2TPv3 pseudowire.

6.3 System Timing Considerations

To ensure proper system operation, the M-CMTS Core SHOULD utilize a timebase that is delayed by no less than 0 and no more than 100 clock cycles (approximately 10 μ s) of the master clock (10.24 MHz) compared to the time communicated via the DTI Client Test Port output.

This time difference between the time communicated via the DTI Client Test Port output and the time reference applied by an M-CMTS Core is expected to be essentially constant. The requirements here do not preclude an M-CMTS Core from taking longer than 100 clock cycles to process the DTI conveyed time, but in that case, the M-CMTS Core would need to internally adjust the applied timebase such that it is delayed between zero and 100 clock cycles from the DTI conveyed timebase.

7 DEPI CONTROL PLANE

Note: This section is normative.

The DEPI control plane is based upon L2TPv3 signaling. The intent of this specification is to follow the provisions of [RFC 3931]. This section includes some examples of how L2TPv3 signaling is used, and includes the extensions and interpretations of the [RFC 3931] specification as it applies to DOCSIS.

All requirements from [RFC 3931] **MUST** be met by both the M-CMTS Core and the EQAM unless this specification explicitly states that a particular requirement from [RFC 3931] is not required.

7.1 Topology




Figure 7-1 - L2TP Topology for Modular CMTS

Figure 7-1 shows how the Modular CMTS architecture maps to L2TP topology. In L2TPv3 nomenclature, the M-CMTS Core and the EQAM are known as a L2TP Access Concentrator (LAC). The M-CMTS Core and the EQAM are considered as peers and may also be referred to as L2TP nodes or as L2TP Control Connection Endpoints (LCCE). For the purposes defined in this document, each LCCE is identified on the network by a single IP address. The connections between two LCCEs are known as Pseudo-Wires (PW).

The L2TP supports both a data path and an inband control path. In L2TPv3 nomenclature, data messages are sent on the data channel, and control messages are sent on the control connection.

First a control connection is established between the two LCCEs, then a session is established. An L2TP session is established before L2TP begins to forward session frames for data. Multiple sessions may be bound to a single control connection.

7.2 Addressing

The M-CMTS Core **SHOULD** use the IP Address of the EQAM and the TSID of the QAM Channel to uniquely identify a QAM Channel within an EQAM during initial configuration.

The M-CMTS Core **MUST** establish at most one control connection for each LCCE pair. This control connection will manage all sessions between the M-CMTS Core and the EQAM. If the UDP Header is used, the control connection **SHOULD** use UDP port 1701.




Figure 7-2 - DEPI Addressing Hierarchy

The M-CMTS Core **MUST** support the creation of a single session per QAM channel. The EQAM **MUST** support a single session per QAM channel. Each DEPI session uses one of the Pseudowire types described in Section 7.5.1.1. Per [RFC 3931], both the M-CMTS Core and the EQAM assign a unique L2TPv3 session ID for each session. The M-CMTS Core **MUST NOT** attempt to create a session to a QAM channel for which the M-CMTS Core already has an active session. Unless specifically configured otherwise, the EQAM **MUST** reject an attempt to establish a session to a QAM channel for which a session has already been established.

The M-CMTS Core **MAY** create multiple PSP flows per session. Different EQAM implementations may support different numbers of PSP flows on any given DEPI session, which is described further in Section 6.1. During L2TPv3 session setup, the EQAM **MUST** assign each flow a unique flow ID. Reassembly, if applicable, is done per flow ID at the EQAM.

L2TPv3 permits the optional use of a UDP Port. L2TPv3 uses the same UDP port number for both the control connection and the data sessions. DEPI also permits the use of UDP ports, primarily for legacy applications. When UDP ports are used, DEPI differs from L2TPv3 in that DEPI permits the data sessions to have different UDP ports than the control connection. Further, DEPI allows each flow within a DEPI session to have a unique UDP port.

The EQAM **MAY** assign each flow a unique UDP Destination Port. The M-CMTS Core **MUST** address DEPI packets using the UDP Destination Port (if the UDP header is used), L2TPv3 Session ID, and flow ID assigned by the EQAM. This is shown in Figure 7-2 and in more detail in Figure 7-3.




Figure 7-3 - DEPI Addressing Hierarchy

7.3 Control Message Format

The format of a DEPI Control Packet, as shown in Figure 7-4 and Figure 7-5, is based on [RFC 3931], with extensions for DOCSIS. The fields which are common with the DEPI Data Packet are described in Section 8.1. Fields which are used differently or are new are described below. Unless otherwise indicated, all values are shown in decimal notation.

The choice of using or not using a UDP header is done through system configuration and is not a negotiated DEPI parameter. The UDP version of DEPI is intended for systems which use the UDP port to map flows to a QAM Channel within an EQAM. The non-UDP version of DEPI is intended for systems which use the L2TPv3 Session ID to map flows to a QAM Channel within an EQAM.

The M-CMTS Core **MUST** support DEPI directly over IP. The M-CMTS **MAY** support DEPI over UDP. The EQAM **MUST** support DEPI directly over IP. The EQAM **MAY** support DEPI over UDP.

7.3.1 Control Message with a UDP Header




Figure 7-4 - DEPI Control Packet with UDP

7.3.2 Control Message without a UDP Header




Figure 7-5 - DEPI Control Packet without UDP

7.3.3 Common Headers for Control and Data Messages

7.3.3.1 Ethernet 802.3 Header

The Ethernet header is defined by [IEEE-802.3]. The Ethernet Destination Address is an individual address. Operation with DEPI over Ethernet group addresses is not specified at this time. The Ethernet Destination Address may be locally or globally administered.

Upon transmission of this frame by the M-CMTS Core, the Ethernet Destination Address will be the Ethernet address of the EQAM or of the next hop router. Upon reception of this frame by the EQAM, the Ethernet Source Address will be the Ethernet address of the output port of the M-CMTS Core or of the previous hop router.

If the networking interface is Ethernet, the M-CMTS Core MUST support the Ethernet header. If the networking interface is Ethernet, the EQAM MUST support the Ethernet header. If another physical layer interface is used instead of Ethernet, then the Ethernet headers would be replaced with the header format pertaining to that physical layer.

7.3.3.2 Ethernet 802.1Q Header

The Ethernet 802.1Q header is defined by [IEEE-802.1Q]. The use of this header is optional and provides frame prioritization and VLAN support at layer 2. The M-CMTS Core **MUST** be able to support the Ethernet 802.1Q header. The M-CMTS Core **SHOULD** support mapping of 802.1Q user priority of tunneled packets based on the DSCP value of a tunnel's IP packet. The EQAM **SHOULD** support the Ethernet 802.1Q header.

7.3.3.3 IPv4 Header

The IP header is defined by [RFC 791]. The IP Source Address is an IP address belonging to the M-CMTS Core. Currently, the IP Destination Address is unicast and is an IP Address belonging to the EQAM. Operation with DEPI over IP Multicast is not specified at this time.

For implementation considerations and for coexistence with network policies that are not amenable to IP fragmentation, EQAMs are not required to perform IP reassembly. The M-CMTS Core **MUST NOT** use IP fragmentation. The M-CMTS Core **MUST** set the IP DF (don't fragment) bit.

The M-CMTS Core **MUST** support a configurable 6 bit Differentiated Service Code Point (DSCP) as defined by [RFC 2983] and [RFC 3260].

The M-CMTS Core **MUST** support the IPv4 header. The EQAM **MUST** support the IPv4 header.

7.3.3.4 UDP Header

The UDP header is defined by [RFC 768]. The value of the UDP Source Port and the UDP Destination Port is determined through the L2TPv3 control plane between the M-CMTS Core and the EQAM. This value **SHOULD** conform to [IANA-PORTS].

When L2TPv3 Control Message Authentication not in use and the UDP header is in use, both the EQAM and M-CMTS Core **MUST** support the generation of UDP checksums as defined in [RFC 768]. The sender **MUST NOT** set the UDP checksum to 0 for L2TPv3 control messages. The receiver **MUST** support validation of the UDP checksum field, as per [RFC 768].

When L2TPv3 Control Message Authentication is in use, the UDP checksum verification is redundant. Therefore, when L2TPv3 Control Message Authentication is in use and the UDP header is in use, the M-CMTS Core and EQAM **MAY** set the UDP checksum field to 0 on control packets when sending a packet, and M-CMTS Core and EQAM **MUST** ignore the UDP checksum field upon receiving a control packet with a UDP checksum of 0.

The M-CMTS Core **MUST** support the UDP header. The EQAM **MUST** support the UDP header.

7.3.3.5 CRC

The CRC is CRC-32 and is defined by [IEEE-802.3].

The M-CMTS Core **MUST** support the CRC field. The EQAM **MUST** support the CRC field.

7.3.4 Specific Headers for Control Messages

7.3.4.1 L2TPv3 Control Header

These fields are defined in [RFC 3931] and are repeated here for reference. The fields have the following significance:

- T** Type bit. The T bit **MUST** be set to 1, indicating that this is a control message.
- L** Length bit. The L bit **MUST** be set to 1, indicating that the Length field is present.
- S** Sequence bit. The S bit **MUST** be set to 1 indicating that sequence numbers (Ns and Nr) are present.
- X** Reserved bits. All reserved bits **MUST** be set to 0 on outgoing messages and ignored on incoming messages.
- Ver** Version. 4 bits. Set to 3.

Length	2 bytes. The Length field indicates the total length of the message in octets, always calculated from the start of the control message header itself, beginning with the T bit. It does not include the Session ID (shown in Figure 7–5) when present.
CCID	Control Connection Identifier. 4 bytes. Negotiated per control connection.
Ns	Sending Sequence Number. 2 bytes. Indicates sending sequence of this control message.
Nr	Received Sequence Number. 2 bytes. Indicates next expected received sequence number.

7.3.4.2 Attribute Value Pairs (AVP)

There can be one or more attribute value pairs (AVPs) per DEPI control message. The fields have the following significance:

M	Mandatory bit. If this bit is set to a 1 and this AVP is rejected, the control connection or session setup in which the AVP is carried will be torn down, according to the procedures outlined in [RFC-L2TPv3].
H	Hidden bit. This bit is set to a 1 when the contents of the AVP message are encrypted and a 0 when the contents of the AVP message are not encrypted. Encryption of AVP messages is not required for DEPI.
Resv	Reserved. 4 bits. Set to all zeros on transmit. Ignored on receive.
Length	10 bits. Equal to the length of the attribute value field plus 6 bytes.
Vendor ID	2 bytes. For AVPs defined by [RFC 3931], this field is set to 0. For AVPs defined by this specification, this field is set to the IANA assigned CableLabs vendor ID of 4491 (0x118B). For AVPs defined outside the scope of this specification, this may be set to a vendor specific ID.
Attribute Type	2 bytes.
Attribute Value	N bytes.
Reserved	8 bits. If an AVP has a reserved field, the bits in this field should be set to 0 on transmit and ignored on the receive.

If an LCCE receives an AVP with a Vendor ID that it doesn't recognize, and the M bit is set to 0, the LCCE MUST silently discard the AVP.

7.4 Signaling

The supported L2TPv3 messages for the DEPI control plane are shown in Figure 7–1:

Table 7–1 - DEPI Control Messages

#	Mnemonic	Name
Control Connection Management		
1	SCCRQ	Start-Control-Connection-Request
2	SCCRP	Start-Control-Connection-Reply
3	SCCCN	Start-Control-Connection-Connected
4	StopCCN	Stop-Control-Connection-Notification
6	HELLO	Hello
20	ACK	Explicit Acknowledgement
Session Management		
10	ICRQ	Incoming-Call-Request
11	ICRP	Incoming-Call-Reply
12	ICCN	Incoming-Call-Connected
14	CDN	Call-Disconnect-Notify
16	SLI	Set Link Info

L2TPv3 outgoing call messages (OCRQ, OCRP, OCCN) and the WAN-Error-Notify (WEN) message are not required to be supported.

There is a reliable control message delivery mechanism that is accomplished either by sending an Explicit Acknowledgement (ACK) message after any of the control messages or by piggybacking an acknowledgment with the Nr and Ns fields in a later control message. If Control Messages are not acknowledged within the time Control Message Timeout (refer to Annex B), then the Control Message MUST be retransmitted up to the value of Control Message Retry Count (refer to Annex B). For example, the Control Message must be retransmitted a total of 10 times using an exponential back-off value starting at 1 second and growing to a maximum value of 8 seconds.

Note: There will be 7 intervals of 8 seconds in this scheme.

Control message authentication MAY be supported. If Control Message Authentication is supported, the methods described in [RFC 3931], section 5.4.1, should be followed.

The following flow diagrams show typical DEPI message exchanges along with the required AVPs from L2TPv3 and DEPI. Optional AVPs are not shown but may also be present.

7.4.1 Control Connection Signaling

7.4.1.1 Control Connection Setup




Figure 7–6 - DEPI Control Connection Setup

In order to tunnel DOCSIS frames over IP using L2TPv3, an L2TPv3 Control Connection is first established as described in [RFC 3931]. Establishment of the control connection involves an exchange of AVPs that identifies the peer and its capabilities. Each control connection has a Control Connection ID which is assigned by the recipient, and negotiated with Connection ID AVPs during the creation of a control connection.

The M-CMTS Core **MUST** support the ability to initiate control connection signaling (L2TPv3 caller). The EQAM **MUST** support the ability to receive incoming control connection requests from the M-CMTS Core (L2TPv3 callee). Establishment of control connections by the EQAM is not required for DEPI and is outside the scope of this specification.

7.4.1.2 Control Connection Teardown




Figure 7-7 - DEPI Control Connection Teardown

Control connection teardown may be initiated by either LCCE and is accomplished by sending a single StopCCN control message. An implementation may shut down an entire control connection and all sessions associated with the control connection by sending the StopCCN. Thus, it is not necessary to clear each session individually when tearing down the whole control connection. The peer receiving the StopCCN message **MUST** maintain session and control state for a period of time equal to StopCCN Timeout (Annex B) after acknowledging the StopCCN. This provision is for managing lost acknowledgments.

In some cases, such as after a reboot, an LCCE may receive a Hello message for a Control Connection for which it is unaware. If the LCCE which received the Hello message with an unknown Connection ID, the LCCE **SHOULD** send a StopCCN message with the Connection ID set to 0, Result Code 2 - General error, and Error Code 1 - No control connection exists yet for this pair of LCCEs. If the LCCE receives a StopCCN message with a Control Connection ID set to 0, it **MUST** tear down all Control Connections with the other LCCE sending the StopCCN.

7.4.1.3 Control Connection Keep-Alive




Figure 7-8 - DEPI Keep-alive

A periodic keep-alive for the control connection is implemented by sending a Hello message if a period of time known as the Hello Timeout (refer to Annex A) has passed without receiving any message (data or control) from the peer.

7.4.2 Session Signaling

7.4.2.1 Session Setup




Figure 7-9 - DEPI Session Setup

After successful control connection establishment, individual sessions may be created. Each session corresponds to a single data stream between the two LCCEs. In addition to the mandatory and optional AVPs in [RFC 3931], the following DEPI specific AVPs are used as part of the session setup.

ICRQ contains the Remote End ID AVP which contains the TSID of the QAM Channel for which the session is intended.

ICRP contains the Remote Session AVP which indicates the session ID that the EQAM wants to use. ICRP also contains a suite of QAM Channel AVPs (refer to Section 7.5.2) which indicate the EQAM's current configuration, which parameters can be changed, EQAM capabilities, and assigned values such as the UDP port values. If these values are not acceptable to the M-CMTS Core, the M-CMTS Core will return a CDN message with the appropriate error code.

ICCN contains the parameters that the M-CMTS Core wants to change. If these parameters are acceptable to the EQAM, the EQAM will return an ACK (either explicit or implicit). If these parameters are not acceptable to the EQAM, the EQAM will return a CDN message with the appropriate error code.

The receipt and processing of the ICCN triggers the initiation of data forwarding within the EQAM for the session. The EQAM **MUST NOT** transmit data on the QAM channel until the session is configured according to the parameters in the ICCN message. The EQAM **SHOULD NOT** buffer data while the session is being configured.

Since the ICCN message may contain AVP parameters that could reconfigure the QAM Channel and disrupt data transmission, the EQAM **SHOULD** set the Circuit Status AVP in the ICRP message to the down state. After the ICCN is received and when the QAM Channel is operational, then the EQAM **SHOULD** send an SLI message to the M-CMTS Core with the Circuit Status AVP set to the up state (refer to 7.5.1.5).

The M-CMTS Core **MAY** also set its circuit status to the down state in the ICRQ message, and then set its circuit status to the up state in the ICCN message or a follow-on SLI message.

The M-CMTS Core **MUST** support the ability to generate session establishment signaling. The EQAM **MUST** support the ability to receive incoming session establishment requests from the M-CMTS Core. Establishment of L2TP sessions by the EQAM is outside the scope of this specification.

7.4.2.2 Session Teardown




Figure 7-10 - DEPI Session Teardown

Session teardown may be initiated by either LCCE and is accomplished by sending a single CDN control message. An implementation may shut down an entire control connection and all sessions associated with the control

connection by sending the StopCCN. Thus, it is not necessary to clear each session individually when tearing down the whole control connection.

7.4.2.3 Session Updates




Figure 7-11 - DEPI Session Updates

If there is a configuration change to one of the EQAM parameters which is described by a DEPI AVP, the M-CMTS Core MUST send the updated AVP to the EQAM with the Set-Link-Info (SLI) message. If there is a configuration change to one of the EQAM parameters which is described by a DEPI AVP, the EQAM MUST send the updated AVP to the M-CMTS Core using an SLI message.

7.4.3 Required and Optional AVPs

In addition to the mandatory and optional AVPs listed in [RFC 3931] and modified by Table 7-3, the following DEPI AVPs listed in Table 7-2 MUST be present in the DEPI control message if mandatory, and MAY be present in DEPI control messages if optional.