

Data-Over-Cable Service Interface Specifications

DCA - MHA v2

Remote Downstream External PHY Interface Specification

CM-SP-R-DEPI-I07-170524

ISSUED

Notice

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

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

© Cable Television Laboratories, Inc. 2014-2017

DISCLAIMER

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

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

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

Document Status Sheet

Document Control Number:	CM-SP-R-DEPI-I07-170524			
Document Title:	Remote Downstream External PHY Interface Specification			
Revision History:	I01 - Released 06/15/2015 I02 - Released 10/01/2015 I03 - Released 01/21/2016 I04 - Released 05/12/2016 I05 - Released 09/23/2016 I06 - Released 01/11/2017 I07 - Released 05/24/2017			
Date:	May 24, 2017			
Status:	Work in Progress	Draft	Issued	Closed
Distribution Restrictions:	Author Only	GL/Member	GL/Member/ Vendor	Public

Key to Document Status Codes

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

Trademarks

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

Table of Contents

1	SCOPE	9
1.1	Introduction and Purpose	9
1.2	MHA v2 Interface Documents	9
1.3	Requirements and Conventions	9
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	OVERVIEW	19
5.1	System Architecture	19
5.1.1	Reference Architecture	19
5.1.2	DEPI Operation	20
5.1.3	Bonding Services Model	20
5.2	Multiple Services Model	21
5.3	R-PHY System for MPEG-TS Video	21
5.3.1	CCAP Core Video Functions	21
5.3.2	RPD Functions	21
5.4	R-PHY System for OOB	22
5.5	Multicast Support on R-PHY	23
6	DEPI ARCHITECTURE	24
6.1	DEPI Data Path	24
6.1.1	DOCSIS D-MPT Data Path	25
6.1.2	PSP Data Path	25
6.1.3	DOCSIS SYNC Message	26
6.1.4	Latency and Skew Requirements for DOCSIS Channels	28
6.2	System Timing Considerations	28
6.3	MPEG-TS Requirements	29
6.4	Multicast DEPI Requirements	29
7	REMOTE PHY CONTROL PLANE PROTOCOL (NORMATIVE)	30
7.1	Remote PHY Topology	30
7.2	Remote PHY Addressing	32
7.3	CCAP Core L2TP Control Message Format	33
7.3.1	L2TP Control Message without UDP Header	33
7.3.2	L2TP Control Message with a UDP Header	34
7.3.3	Common Headers for Control and Data Messages	35
7.3.4	Specific Headers for Control Messages	40
7.4	L2TP Signaling	41
7.4.1	Control Connection Signaling	41
7.4.2	Session Signaling	44
7.4.3	Mandatory and Optional AVPs	47
7.5	AVP Definition Updates	48
7.5.1	Conventional L2TPv3 AVPs	48
7.5.2	DEPI Control Specific AVP	55
7.5.3	DEPI Session Specific AVPs	56
8	DEPI FORWARDING PLANE	67

8.1	L2TPv3 Transport Packet Format	67
8.1.1	<i>Data Message Format without UDP</i>	67
8.1.2	<i>Data Message with a UDP Header</i>	68
8.1.3	<i>Specific Headers for Data Messages</i>	68
8.2	DEPI MPT Sublayer Header	69
8.3	DEPI Multi-channel MPEG (MCM) L2TPv3 Sublayer Header Format	70
8.4	PSP Sublayer Header	72
8.4.1	<i>PSP Legacy Sublayer Header</i>	72
8.4.2	<i>PSP Sublayer Header</i>	73
8.4.3	<i>Interpretation of PSP Sublayer Header for PLC Sub-channel</i>	75
8.5	Optional Headers	77
8.6	DEPI Latency Measurement (DLM) Sublayer Header	78
8.6.1	<i>DEPI Latency Measurement</i>	79
8.6.2	<i>DEPI Buffer Depth Alert Message</i>	80
8.7	DEPI Multicast Packet Header Format	81
8.8	CCAP Core Output Rate	81
8.8.1	<i>Downstream Output Buffer History</i>	82
8.8.2	<i>Downstream Buffer Depth Monitoring Alerts</i>	83
8.9	Bidirectional Forwarding Detection	83
8.10	IPv6 Support	85
ANNEX A	PARAMETERS AND CONSTANTS (NORMATIVE)	86
ANNEX B	DOCS-IF-M-CMTS-MIB (NORMATIVE)	87
ANNEX C	FORMAT AND CONTENT FOR EVENT, SYSLOG, AND SNMP NOTIFICATION (NORMATIVE)	88
ANNEX D	DEPI PATH REDUNDANCY (NORMATIVE)	89
APPENDIX I	DEPI AND DOCSIS SYSTEM PERFORMANCE (INFORMATIVE)	90
APPENDIX II	ACKNOWLEDGEMENTS	91
APPENDIX III	REVISION HISTORY	92

List of Figures

Figure 1 - MHA v2 / R-PHY System Architecture	19
Figure 2 - Downstream Channel Bonding Services Model	21
Figure 3 - Block Diagram of Typical Remote PHY Setup	22
Figure 4 - Downstream DOCSIS and Video R-PHY Block Diagram	24
Figure 5 - Format of a DOCSIS SYNC MAC Message	26
Figure 6 - End-to-end Remote PHY Topology	30
Figure 7 - Remote PHY Topology with Separate CMTS Core and EQAM Cores	31
Figure 8 - Remote PHY Topology with Multiple Connections	31
Figure 9 - R-PHY Connectivity Hierarchy	32
Figure 10 - L2TP Control Message without UDP Header	34
Figure 11 - L2TPv3 Control Packet with UDP Header	35
Figure 12 - R-PHY L2TP Control Connection Setup	42
Figure 13 - R-PHY L2TP Control Connection Teardown	43
Figure 14 - R-PHY L2TP Keep-alive	44
Figure 15 - R-PHY L2TP Session Setup	44
Figure 16 - R-PHY Session Teardown	46

Figure 17 - R-PHY Session Updates47

Figure 18 - Message Type AVP.....49

Figure 19 - Result Code AVP49

Figure 20 - Host Name AVP.....49

Figure 21 - Vendor Name AVP50

Figure 22 - Serial Number AVP50

Figure 23 - Control Connection DSCP AVP50

Figure 24 - Router ID AVP.....50

Figure 25 - Assigned Control Connection ID AVP51

Figure 26 - Pseudowire Capabilities List AVP51

Figure 27 - Local Session ID AVP51

Figure 28 - Remote Session ID AVP52

Figure 29 - Remote End ID AVP.....52

Figure 30 - RPD Channel Selector.....52

Figure 31 - Pseudowire Type AVP53

Figure 32 - L2-Specific Sublayer AVP.....53

Figure 33 - Data Sequencing AVP.....54

Figure 34 - Circuit Status AVP54

Figure 35 - DEPI Pseudowire Subtype Capabilities List AVP55

Figure 36 - DEPI Result and Error Code AVP57

Figure 37 - DEPI Resource Allocation Request AVP.....57

Figure 38 - DEPI Resource Allocation Reply AVP59

Figure 39 - DEPI Local MTU AVP59

Figure 40 - DEPI Remote MTU Max Payload AVP.....60

Figure 41 - DEPI Local UDP Port AVP60

Figure 42 - DEPI Remote Multicast Join AVP.....61

Figure 43 - DEPI Remote Multicast Leave AVP.....61

Figure 44 - DEPI Multicast Capability AVP62

Figure 45 - Upstream Flow AVP63

Figure 46 - DEPI Pseudowire Subtype AVP63

Figure 47 - DEPI L2-Specific Sublayer Subtype AVP64

Figure 48 - DEPI DLM Response DSCP AVP.....64

Figure 49 - DEPI S-BFD Return Path DSCP AVP.....65

Figure 50 - DEPI Multicast UDP Port AVP65

Figure 51 - L2TPv3 Data Packet Outer Encapsulation (without UDP)67

Figure 52 - L2TPv3 Data Packet Outer Encapsulation with UDP Header.....68

Figure 53 - DEPI MPT Sublayer Header and Payload69

Figure 54 - MCM Sublayer Header and Payload.....71

Figure 55 - DEPI Legacy PSP Sublayer Header and Payload72

Figure 56 - DEPI PSP Sublayer Header and Payload74

Figure 57 - Use of DEPI PSP Sublayer Header for PLC Sub-channel76

Figure 58 - Example of a Complete DEPI Packet Format (IPv4) with Extras.....78

Figure 59 - DEPI Latency Measurement (DLM) Sublayer Header for Latency Measurements.....79

Figure 60 - DLM Sub-layer Header for Buffer Depth Monitor Alert Function.....80

Figure 61 - BFD Sublayer (with UDP).....	84
--	----

List of Tables

Table 1 - List of MHA v2 Specifications.....	9
Table 2 - DOCSIS SYNC MAC Message Fields.....	27
Table 3 - L2TPv3 Control Header Fields.....	40
Table 4 - R-PHY Control Message AVPs	40
Table 5 - R-PHY L2TP Control Messages	41
Table 6 - R-PHY L2TP Mandatory and Optional AVPs[16].....	47
Table 7 - R-PHY Supported L2TPv3 AVPs.....	48
Table 8 - Pseudowire Types.....	51
Table 9 - RPD Channel Selector.....	53
Table 10 - L2-Specific Sublayer Types	54
Table 11 - Circuit Status AVP Details.....	54
Table 12 - DEPI Pseudowire Subtypes.....	55
Table 13 - DEPI Defined Session Specific AVPs.....	56
Table 14 - DEPI Result and Error Code AVP Details	57
Table 15 - DEPI Resource Allocation Request AVP Details.....	58
Table 16 - DEPI Resource Allocation Reply AVP Details.....	59
Table 17 - DEPI Local MTU AVP Details	59
Table 18 - DEPI Remote MTU Max Payload AVP Details.....	60
Table 19 - DEPI Local UDP Port AVP Details	60
Table 20 - DEPI Remote Multicast Join AVP Details	61
Table 21 - DEPI Remote Multicast Leave AVP Details.....	62
Table 22 - DEPI Multicast Capability AVP Details	62
Table 23 - Upstream Flow AVP Details	63
Table 24 - DEPI L2-Specific Sublayer Subtype Values	64
Table 25 - DEPI DLM Response DSCP AVP Details	65
Table 26 - DEPI S-BFD Return Path DSCP AVP Details.....	65
Table 27 - DEPI Multicast UDP Port AVP Details	65
Table 28 - Header Format for MPT Sublayer Header and Payload Details.....	69
Table 29 - Sublayer “H” Bits Decode Table.....	70
Table 30 - L2TPv3 MCM Sublayer Header and Payload Format.....	71
Table 31 - DEPI Legacy PSP Sublayer Header and Payload Format Details.....	73
Table 32 - DEPI PSP Sublayer Header and Payload Format Details.....	74
Table 33 - DEPI Latency Measurement (DLM) Sublayer Header for Latency Measurement Details	79
Table 34 - Details of DLM Sublayer Header for Buffer Depth Monitor Alert Function.....	80
Table 35 - Parameters and Constants.....	86

This page has been left blank intentionally

1 SCOPE

1.1 Introduction and Purpose

This document describes the Remote DEPI and control plane operation of the Modular Headend Architecture version 2 (MHA_{v2}). MHA_{v2}/Remote PHY Architecture permits a CMTS to support an IP based digital HFC plant. In an IP based digital HFC plant, the fiber portion utilizes a baseband network transmission technology such as Ethernet, EPON (Ethernet over Passive Optical Networks), GPON (Gigabit Passive Optical Network), or any Layer 2 technology that would support a fiber based Layer 1. MHA_{v2} uses a Layer 3 pseudowire between a CCAP Core and a series of Remote PHY Devices. One of the common locations for a Remote PHY Device is located at an optical node device that is located at the junction of the fiber and coax plants.

This document was created from the basis of the I08 version of the MHA DEPI specification, [DEPI].

1.2 MHA_{v2} Interface Documents

A list of the documents in the MHA_{v2} family of specifications is provided below. For updates, refer to <http://www.cablelabs.com/specs/specification-search/>.

Table 1 - List of MHA_{v2} Specifications

Designation	Title
CM-SP-R-PHY	Remote PHY Specification
CM-SP-R-DEPI	Remote Downstream External PHY Interface Specification
CM-SP-R-UEPI	Remote Upstream External PHY Interface Specification
CM-SP-GCP	Generic Control Plane Specification
CM-SP-R-DTI	Remote DOCSIS Timing Interface Specification
CM-SP-R-OOB	Remote Out-of-Band Specification
CM-SP-R-OSSI	Remote PHY OSS Interface Specification

NOTE: MHA_{v2} does not explicitly use any of the original Modular Headend Architecture specifications.

1.3 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, and users of this document are encouraged to investigate the possibility of applying the most recent editions of the documents listed below. References are either specific (identified by date of publication, edition number, version number, etc.) or non-specific. For a nonspecific reference, the latest version applies.

[DEPI]	Downstream External PHY Interface Specification, CM-SP-DEPI-I08-100611, June 11, 2010, Cable Television Laboratories, Inc.
[DRFI]	DOCSIS Downstream Radio Frequency Interface, CM-SP-DRFI-I16-170111, January 11, 2017, Cable Television Laboratories, Inc.
[DTI]	DOCSIS Timing Interface, CM-SP-DTI-I06-150305, March 5, 2015, Cable Television Laboratories, Inc.
[EQAM-VSI]	Edge QAM Video Stream Interface Specification, CM-SP-EQAM-VSI-I01-081107, November 7, 2008, Cable Television Laboratories, Inc.
[GCP]	Generic Control Plane Specification, CM-SP-GCP-I03-170524, May 24, 2017, Cable Television Laboratories, Inc.
[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.
[IEEE 802.3as]	802.3as-2006 - IEEE Standard for 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
[ISO 13818-1]	ISO/IEC 13818-1:2013, Information technology, Generic Coding of Moving Pictures and Associated Audio Information. Part 1: System, May 23, 2013.
[ITU-T J.83]	ITU-T Recommendation J.83 (4/97), Digital multi-programme systems for television sound and data services for cable distribution.
[MULPIv3.1]	DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I11-170510, May 10, 2017, Cable Television Laboratories, Inc.
[PHYv3.1]	DOCSIS 3.1 Physical Layer Specification, CM-SP-PHYv3.1-I11-170510, May 10, 2017, Cable Television Laboratories, Inc.
[R-DTI]	Remote DOCSIS Timing Interface Specification, CM-SP-R-DTI-I05-170524, May 24, 2017, Cable Television Laboratories, Inc.
[R-OOB]	Remote Out-of-Band Specification, CM-SP-R-OOB-I06-170524, May 24, 2017, Cable Television Laboratories, Inc.
[R-OSSI]	Remote PHY OSS Interface Specification, CM-SP-R-OSSI-I06-170524, May 24, 2017, Cable Television Laboratories, Inc.
[R-PHY]	Remote PHY Specification, CM-SP-R-PHY-I07-170524, May 24, 2017, Cable Television Laboratories, Inc.
[R-UEPI]	Remote Upstream External PHY Interface Specification, CM-SP-R-UEPI-I06-170524, May 24, 2017, Cable Television Laboratories, Inc.
[RFC 768]	IETF RFC 768, User Datagram Protocol, August 1980.
[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 2460]	IETF RFC 2460, Internet Protocol, Version 6 (IPv6) Specification, December 1998.
[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 3308]	IETF RFC 3308, Layer Two Tunneling Protocol (L2TP) Differentiated Services Extension, November 2002.
[RFC 3931]	IETF RFC 3931, Layer Two Tunneling Protocol - Version 3 (L2TPv3), March 2005.
[RFC 5085]	IETF RFC 5085, Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires, December 2007.
[RFC 6935]	IETF RFC 6935, IPv6 and UDP Checksums for Tunneled Packets, April 2013.
[RFC 7880]	IETF RFC 7880, Seamless Bidirectional Forwarding Detection (S-BFD), July 2016.
[RFC 7885]	IETF RFC 7885, Seamless Bidirectional Forwarding Detection (S-BFD) for Virtual Circuit Connectivity Verification (VCCV), July 2016.
[RFC 7886]	IETF RFC 7886, Advertising Seamless Bidirectional Forwarding Detection (S-BFD) Discriminators in the Layer Two Tunneling Protocol Version 3 (L2TPv3), July 2016.
[RFIv2.0]	DOCSIS Radio Frequency Interface Specification, CM-SP-RFIv2.0-C02-090422, April 22, 2009, Cable Television Laboratories, Inc.
[RMI-ERM-EDGE]	Edge Resource Manager - Edge Device Interface Specification, CM-SP-RMI-ERM-EDGE-I02-150528, May 28, 2015, Cable Television Laboratories, Inc.
[S-BFD-BASE]	S-BFD Base Internet draft, http://datatracker.ietf.org/doc/draft-ietf-bfd-seamless-base
[S-BFD-VCCV]	S-BFD for VCCV, http://tools.ietf.org/html/draft-gp-pals-seamless-vcv-00
[S-BFD-DSR-L2TP]	S-BFD discriminator for L2TPv3, http://tools.ietf.org/html/draft-gp-l2tpext-sbfd-discriminator-00

2.2 Informative References

This document uses the following informative references:

[IANA-L2TP]	IANA, Layer Two Tunneling Protocol (L2TP) Parameters.
[ISO 8802-2]	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.
[RFC 5580]	IETF RFC 5580, Bidirectional Forwarding Detection (BFD), June 2010.
[SCTE 55-1]	ANSI/SCTE 55-1 2009, Digital Broadband Delivery System: Out of Band Transport Part 1: Mode A.
[SCTE 55-2]	ANSI/SCTE 55-2 2008, Digital Broadband Delivery System: Out of Band Transport Part 2: Mode B.

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>
- European Telecommunications Standards Institute, ETSI, <http://www.etsi.org>
- The Institute of Electrical and Electronics Engineers, Inc., IEEE, Internet: <http://standards.ieee.org>
- Internet Assigned Numbers Authority, IANA, Internet: <http://www.iana.org>
- Internet Engineering Task Force (IETF) Secretariat, 48377 Fremont Blvd., Suite 117, Fremont, California 94538, USA, Phone +1-510-492-4080, Fax: +1-510-492-4001. <http://www.ietf.org>.
- International Organization for Standardization (ISO), Tel.: +41 22 749 02 22, Fax: +41 22 749 01 55, www.standardsinfo.net
- International Telecommunication Union (ITU), Phone +41 22 730 5111 (ITU Switchboard), Fax +41 22 733 7256, <http://www.itu.int/>

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.
CCAP Core	A CCAP device that uses MHA v2 protocols to interconnect to an RPD.
Converged Interconnect Network	The network (generally gigabit Ethernet) that connects a CCAP Core to an RPD.
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 $dB = 10 \log_{10}(P_{OUT}/P_{IN})$.
Decibel-Millivolt (dBmV)	Unit of RF power expressed in decibels relative to 1 millivolt, where $dBmV = 20 \log_{10}(\text{value in mV}/1 \text{ mV})$.
Downstream (DS)	<ol style="list-style-type: none"> 1. Transmissions from CMTS to CM. This includes transmission from the CCAP Core to the RPD, as well as the RF transmissions from the RPD 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).
Ephemeral	A TCP or UDP port number that is not assigned by the Internet Assigned Numbers Authority (IANA).
Flow	A stream of packets in DEPI used to transport data of a certain priority from the CCAP Core to a particular QAM channel of the RPD. 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 an 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
Packet Identifier (PID)	A unique integer value used to identify elementary streams of a program in a single or multi-program Transport Stream as described in section 2.4.3 of ITU-T Rec. H.222.0 [ISO 13818-1].
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.
Program Clock Reference	A timestamp in the Video Transport Stream from which decoder timing is derived.
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/ .

Remote PHY Device	The RPD is a device in the network which implements the Remote PHY specification to provide conversion from digital Ethernet transport to analog RF transport.
Session	An L2TPv3 data plane connection from the CCAP Core to the QAM channel. There must be one session per Downstream 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)	<ol style="list-style-type: none">1. Transmissions from CM to CMTS. This includes transmission from the RPD to CCAP Core as well as the RF transmissions from the CM to the RPD.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
ATM	Asynchronous Transfer Mode
AVP	L2TPv3 Attribute Value Pair
BAT	Buffer Alert Threshold
BFD	Bidirectional Forwarding Detection
CCAP™	Converged Cable Access Platform
CDN	L2TPv3 Call-Disconnect-Notify message
CIN	Converged Interconnect Network
CM	Cable Modem
CMCI	Cable Modem to Customer Premise Equipment Interface,
CMTS	Cable Modem Termination System
CPE	Customer Premises Equipment
CRC	Cyclic Redundancy Check
CRC16	CRC of length 16
CSMA	Carrier Sense Multiple Access
CV	Connectivity Verification
DA	Destination Address
dB	Decibels
dBmV	Decibel-Millivolt
DCA	Distributed CCAP Architecture
DEPI	Downstream External-PHY Interface
DLM	DEPI Latency Measurement
DOCSIS	Data-Over-Cable Service Interface Specifications
DOCSIS-MPT (D-MPT)	DOCSIS MPT Mode
DRFI	Downstream Radio Frequency Interface
DS	Downstream
DSAP	Destination Service Access Point
DSCP	Differentiated Services Code Point
DTI	DOCSIS Timing Interface
DTS	DOCSIS Time Stamp, 32-bit
EMA	Exponential Moving Average
EQAM	Edge QAM
ERM	Edge Resource Manager
ERMI	Edge Resource Manager Interface
FC	Frame Control
FEC	Forward Error Correction
FIFO	First In, First Out
FQDN	Fully Qualified Domain Name
Gbps	Gigabits per second
GCP	Generic Control Plane
GE	Gigabit Ethernet (Gig E)
GHz	Gigahertz
HCS	Header Check Sequence
HDLC	High-Level Data Link Control

HELLO	L2TPv3 Hello message
HFC	Hybrid Fiber/Coax
Hz	Hertz
ICCN	L2TPv3 Incoming-Call-Connected message
ICRP	L2TPv3 Incoming-Call-Reply message
ICRQ	L2TPv3 Incoming-Call-Request message
ID	Identifier
IEEE	Institute of Electrical and Electronic Engineers
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IKE	Internet Key Exchange
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
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
LAG	Link Aggregation Group
LCCE	L2TP Control Connection Endpoint
LNS	L2TP Network Server
LSB	Least Significant Bit
MAC	Media Access Control
MAP	Upstream Bandwidth Allocation Map (referred to only as MAP)
Mbps	Megabits per second
MCM	Multi Channel MPEG
M-CMTS	Modular Cable Modem Termination System
MER	Modulation Error Ratio
MHz	Megahertz
MIB	Management Information Base
MLD	Multicast Listener Discovery
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
MPLS	Multiprotocol Label Switching
MPT	MPEG-TS mode of DEPI
MPTS	Multi Program Transport Stream
msec	Millisecond
MSO	Multiple System Operator
MSB	Most Significant Bit
MTU	Maximum Transmission Unit
NCO	Numerically Controlled Oscillator
NDF	Narrowband Digital Forward

NDR	Narrowband Digital Reverse
NF	Normalization Factor
ns	Nanosecond
NSI	Network Side Interface
OFDM	Orthogonal Frequency Division Multiplexing
OFDMA	Orthogonal Frequency Division Multiple Access
OOB	Out-of-band
OSSI	Operations System Support Interface
PCR	Program Clock Reference
PDU	Protocol Data Unit
PHB	Per Hop Behavior
PIM	Protocol-Independent Multicast
PHY	Physical Layer
PID	Packet Identifier
PHB-ID	Per Hop Behavior Identifier
PLC	PHY Link Channel
PMD	Physical Media Dependent Sublayer
PPP	Point-to-Point Protocol
PSI	Program Specific Information
PSP	Packet-Streaming-Protocol
PW	Pseudowire
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
R-DEPI	Remote DEPI
R-DTI	Remote DTI
REQ	Request
RF	Radio Frequency
RFI	Radio Frequency Interface
RFC	Request For Comments
RPD	Remote PHY Device
R-PHY	Remote PHY
R-OOB	Remote Out-of-band
R-UEPI	Remote UEPI
SA	Source Address
SAR	Segmentation and Reassembly
S-BFD	Seamless BFD
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
SC-QAM	Single Carrier Quadrature Amplitude Modulation
SDV	Switched Digital Video
SID	Service ID
SLA	Service Level Agreement
SLI	L2TPv3 Set Link Info message
SPR	Stream Processing Resource
SPTS	Single Program Transport Stream
SSAP	Source Service Access Point
StopCCN	L2TPv3 Stop-Control-Connection-Notification message

TC	Transmission Convergence
TCP	Transmission Control Protocol
TS	Transport Stream
TSID	MPEG2 Transport Stream ID
UCD	Upstream Channel Descriptor
UDP	User Datagram Protocol
UEPI	Upstream External-PHY Interface
UML	Unified Modeling Language
US	Upstream
VCCV	Virtual Circuit Connectivity Verification
VoD	Video On Demand
WFQ	Weighted Fair Queuing

5 OVERVIEW

5.1 System Architecture

5.1.1 Reference Architecture

The architecture for an MHA_{v2} system is shown in Figure 1. This architecture contains several pieces of equipment along with interfaces between those pieces of equipment. This section briefly introduces each device and interface.




Figure 1 - MHA_{v2} / R-PHY System Architecture

The Remote PHY Device (RPD) is a device that has a network interface on one side and an RF interface on the other side. The RPD provides Layer 1 PHY conversion, Layer 2 MAC conversion, and Layer 3 pseudowire support. The RPD RF output may be combined with other overlay services such as analog or digital video services.

The CCAP Core contains everything a traditional CCAP does, except for functions performed in the RPD. The CCAP Core contains the downstream DOCSIS MAC, the upstream DOCSIS MAC, all the initialization and operational DOCSIS-related software as well as the majority of the video EQAM functions.

Note that the original MHA architecture had the downstream PHY external and the upstream PHY internal. MHA was used to interface to an EQAM (Edge QAM) device that was co-located at the head end with the CMTS Core. Thus, the main differences between MHA_{v1} and MHA_{v2} are the location of the upstream PHY, the division of EQAM functions between the CCAP Core and the RPD, as well as the role of the solution in the marketplace. From a technical standpoint, the solutions are very similar.

Due to the physical separation of the downstream PHY and the upstream PHY in MHA_{v1}, a DOCSIS Timing Interface (DTI) Server is needed to provide a common frequency of 10.24 MHz and a DOCSIS timestamp between the two MHA_{v1} elements. In MHA_{v2}, the same DTI server is not required because the downstream and upstream PHYs are co-located in the RPD. A different timing solution referred to as **R-DTI** is used to provide timing services for functions such as DOCSIS scheduling.

R-DEPI, the Downstream External PHY Interface, is the downstream interface between the CCAP Core and the RPD. R-DEPI is based on [DEPI]. More specifically, it is an IP pseudowire between the MAC and PHY in an MHA_{v2} system that contains both a data path for DOCSIS frames, video packets, and OOB packets, as well as a control path for setting up, maintaining, and tearing down sessions. MHA_{v1} used the MPT (MPEG-TS) encapsulation. MHA_{v2} retains the original MPT encapsulation for backward compatibility but also added a new MPEG encapsulation called MCM (Multi-channel MPEG). MHA_{v2} also requires the PSP (Packet Streaming Protocol) mode for transport of DOCSIS data, and supports new services like DOCSIS 3.1.

R-UEPI, the Upstream External PHY Interface, is the upstream interface between the RPD and the CCAP Core. Like DEPI, it is an IP pseudowire between the PHY and MAC in an MHA v2 system that contains both a data path for DOCSIS frames, and a control path for setting up, maintaining, and tearing down sessions.

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 10 Gbps Ethernet.

CMCI, or Cable Modem to Customer Premise Equipment Interface, is also unchanged, and is typically Ethernet, USB, or Wi-Fi.

5.1.2 DEPI Operation

DEPI is an IP Tunnel that exists between the DOCSIS or MPEG Video MAC in the CCAP Core and the PHY that exists in the RPD. 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 RPD for transmission.

The base protocol that is used for the DEPI is the Layer 2 Tunneling Protocol version 3, or L2TPv3 for short (see [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, and PPP.

Section 8.1 shows the format of an L2TPv3 data packet. Each data packet contains a 32-bit session ID that is associated with a single QAM Channel or a group of QAM channels.

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

There are three 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 subheader which contains a sequence number so packet drops can be detected. The encapsulation of DOCSIS frames into MPEG-TS packets is performed in the CCAP Core.

The second technique, known as Multi-Channel-MPEG (or MCM) is a variation of the D-MPT technique. MCM allows a single DEPI packet to transport MPEG frames for multiple QAM channels. D-MPT and MCM can be utilized to carry data to SC-QAM channels.

The third technique, known as the Packet Stream Protocol (PSP) can be used to transport DOCSIS frames and PLC information in the L2TPv3 payload. PSP can be utilized for transport of data to SC-QAM channels and OFDM channels. When sent to a SC-QAM channel, the DOCSIS frames are then encapsulated in MPEG-TS packets within the RPD. When sent to an OFDM channel, the OFDM Codeword Builder converts PSP packets into OFDM codewords and PLC messages. 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 MHA v2 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.

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 RPD.

5.1.3 Bonding Services Model

Downstream Channel Bonding refers to the forwarding of a stream of DOCSIS frames across multiple QAM carriers. The specifics of the forwarding are shown in Figure 2.




Figure 2 - Downstream Channel Bonding Services Model

In the MHAv2 architecture, the Downstream Channel Bonding is implemented in the CCAP Core. At the CCAP Core, a packet from the IP backbone is 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 over DEPI using D-MPT, MCM, or PSP techniques.

In this system, the RPD is unaware that it is performing bonding. It is also unaware of any of the details of the DOCSIS bonding protocol.

5.2 Multiple Services Model

DEPI supports both DOCSIS services and MPEG video services. The DOCSIS services and MPEG video services are delivered on a separate set of downstream channels and typically transported over DEPI through a separate set of pseudowires.

5.3 R-PHY System for MPEG-TS Video

The long-established and widely understood role of the EQAM in the video-on-demand (VoD) and switched-digital-video (SDV) architecture is to receive an IP unicast or multicast stream containing MPEG transport stream packets, and then produce that transport stream on one or more RF outputs for transmission over the hybrid fiber-coax (HFC) cable plant.

The chosen approach for MPEG video support in R-PHY architecture is to have the CCAP Core perform the vast majority of the EQAM functions and have the RPD perform minimal functions. With this in mind, the CCAP Core receives SPTSs or MPTSs, multiplexes the streams, and generates an MPEG-compliant MPTS that is then transported over L2TPv3 to the RPD. The RPD then performs de-jittering to handle the jitter that the intervening network had introduced.

5.3.1 CCAP Core Video Functions

The CCAP Core receives SPTSs or MPTSs and outputs MPTSs. It performs a number of functions such as demultiplexing, PID remapping/filtering, multiplexing, PSI parsing and re-generation, de-jittering, clock recovery and PCR restamping, NULL packet insertion, and encryption as described in [EQAM-VSI]. After performing all these functions, the CCAP Core outputs a constant bit rate multiplex output and transports the multiplexed stream over L2TPv3, with up to 10 MPEG packets in one L2TPv3 packet.

The CCAP Core also participates in the required control plane signaling for setup and teardown of sessions as described in [RMI-ERM-EDGE].

5.3.2 RPD Functions

The RPD receives only constant bit rate MPTSs which themselves may contain streams for VoD/SDV/Broadcast applications. The RPD is unaware of any individual video sessions, or their setup and teardown. The rate of the incoming MPTS is signaled to the RPD from the CCAP Core via the control plane.

There are two possible modes of operation for the RPD, both of which are described below.

5.3.2.1 Asynchronous Mode

In this mode of operation, the RPD does not make any assumptions about the clock frequencies of the CCAP Core and RPD being synchronized. Such an RPD has a MPEG-compliant clock and adds and/or deletes NULL packets to account for the clock frequency difference between the CCAP Core and RPD. A block diagram of a Remote PHY implementation as described above is shown in Figure 3.




Figure 3 - Block Diagram of Typical Remote PHY Setup

The incoming MPTS is de-jittered to handle any jitter that was introduced in the intervening network. The incoming MPTS is assumed to be constant bit rate and the rate is known at the RPD. This allows the Remote PHY to detect if the clock frequency of the CCAP Core is ahead or behind the RPD clock, and hence delete or insert NULL packets as needed. The RPD assumes that the PCR that was stamped in the incoming packets was MPEG compliant and hence only adjusts the PCR as needed to account for any NULL packets introduced or deleted in the RPD. For the system to operate properly, enough NULL packets should be present at the output of the CCAP Core so that the RPD can delete them as needed to adjust for clock frequency differences.

5.3.2.2 Synchronous Mode

In this mode of operation, the RPD clock is synchronized to the CCAP Core clock via R-DTI. An RPD that operates in such a synchronous mode could in fact avoid PCR re-stamping, by simply replacing lost packets with NULL packets. Lost packets could be detected by checking sequence numbers contained in D-MPT headers. However, the RPD will still need to de-jitter the incoming MPTS to remove the effects of the CIN jitter.

The RPD MUST insert NULL MPEG packets to compensate for any video D-MPT frame lost in the CIN, as indicated by the sequence number in the D-MPT header. The number of NULL MPEG packets to insert will equal the number of the MPEG packets contained in the last received D-MPT frame. The RPD MUST insert the NULL MPEG packets at the exact transport stream position of the lost MPEG packets, when no more than one consecutive D-MPT frame is lost.

In the event that more than one consecutive frame is lost, the implementation is vendor specific.

5.4 R-PHY System for OOB

The R-PHY System for Out-of-Band control data is described in the [R-OOB] specification.

5.5 Multicast Support on R-PHY

The primary use case for multicast delivery between CCAP Core and the RPD is for the delivery of broadcast video services from a single CCAP Core element to a number of RPDs. This allows the system to scale by allowing a single CCAP Core element to generate and serve streams to all the RPDs that are configured to receive the same broadcast lineup. Since broadcast Serving Groups are quite large (~100,000 or more subscribers), using multicast to deliver the same copy to hundreds of Remote PHY devices provides significant cost savings for operators.

The same mechanism may also be used for other purposes, such as to replicate a group of narrowcast video channels (e.g., Video-On-Demand or Switched Digital Video) to multiple RPDs, or to replicate DOCSIS channels so as to combine multiple RPDs into a single DOCSIS service group.

This mechanism is meant for the replication of an entire QAM or OFDM channel to multiple RPDs. The model is that the exact contents of the channel are created at the CCAP Core - i.e., the CCAP Core creates the full MPTS multiplex for MPEG-TS video channels, or performs all required functions to schedule IP packets onto the downstream for DOCSIS channels. The data plane between CCAP Core and RPD consists of D-MPT or PSP pseudowires which are identical to those used for unicast delivery to a single RPD, except that they use IP multicast addressing (as specified in later sections of this document).

With this mechanism, the CIN only sees standard IP multicast packets, and hence, standard multicast forwarding works as expected. The RPD participates in multicast signaling using IGMPv3/MLDv2, which ensures proper forwarding of the multicast packets over the CIN. The RPD learns which multicast streams to join via the L2TPv3 control plane. The L2TPv3 control plane itself remains point-to-point.

It is expected that the CCAP Core will use a single layer 2/3 multicast group address to carry one or more D-MPT and/or PSP pseudowires for a number of QAM or OFDM channels, grouping them in accordance with the service groups reached. Multicast session IDs (as defined later in this document) will then be used to differentiate individual pseudowires corresponding to different downstream channels or groups of channels. This addressing model is very similar to that used for unicast delivery.

As an example, the CCAP Core might use a single layer 2/3 (S,G) for all channels in the broadcast video lineup being distributed to a large number of RPDs. The same CCAP Core (or possibly a different one) would then use an additional layer 2/3 (S,G) for narrowcast service groups for VOD or SDV. Each such group might contain on the order of 8-12 QAM channels and reach on the order of 4-8 RPDs. Similar groups could exist for bundles of DOCSIS downstream channels representing a complete or partial Downstream Service Group, with each such group being distributed to 2-4 RPDs.

OOB signals could also be distributed using multicast addressing. [R-OOB] describes the details of formatting for these signals. OOB pseudowires might originate from a different CCAP Core than those used for DOCSIS and video, and will very likely reach a different group of RPDs than other services. Thus, the CCAP Core would use additional multicast layer 2/3 addresses to carry pseudowires for OOB signals.

Using this approach, an RPD supporting a single downstream port might see a configuration including (as an example) 1-2 multicast addresses for broadcast services, 2-4 multicast addresses for narrowcast video services, and 1-2 multicast addresses for groups of DOCSIS downstream channels. An additional 1-2 addresses could be used for OOB (for instance, one for either 55-1 or 55-2, and another for NDF signals). This would give a total of 10 multicast groups to be joined by the RPD per downstream RF port.

6 DEPI ARCHITECTURE

6.1 DEPI Data Path

A simplified logical block diagram of the internal downstream DOCSIS and video data path of the RPD is shown in Figure 4.




Figure 4 - Downstream DOCSIS and Video R-PHY Block Diagram

The requirements for processing MPEG Video are described in Section 6.3.

The CCAP Core MUST support PSP mode for transport of data on DOCSIS SC-QAM and OFDM channels.

The CCAP Core MAY support D-MPT mode or MCM Mode, or both modes for transport of data on DOCSIS SC-QAM channels.

The RPD MUST support PSP mode for transport of data on DOCSIS SC-QAM and OFDM channels.

The RPD MAY support D-MPT mode, MCM Mode, or both modes for transport of data on DOCSIS SC-QAM channels.

The CCAP Core creates an L2TPv3 session with either a PSP or non-PSP pseudowire type (e.g., D-MPT or MCM). The RPD is not intended to simultaneously support PSP and non-PSP traffic within a single session. For non-PSP sessions, the RPD transmits information from L2TPv3 data packets directly to a downstream RF channel in order of packet sequence number. The "FlowId" field of the L2TPv3 sublayer for non-PSP sessions is reserved; it is transmitted as 0 and ignored on receipt. For PSP sessions, packets are grouped into one or more PSP "flows" in each direction. The RPD supports reception and assembly of PSP-fragmented packets on each downstream flow. The RPD enqueues each assembled downstream packet onto a "QoS queue" for scheduling by a Packet Scheduler onto a downstream channel.

The RPD is not intended to simultaneously support D-MPT and PSP within a single session.

For both non-PSP and PSP modes, the RPD MUST insert null MPEG packets when it has no data to send. The RPD 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 (see Section 6.1.3).

In the upstream direction, an RF channel is assigned by the CCAP Core (via GCP) to a set of PSP sessions. Upstream data PSP sessions may have more than one PSP flow; other pseudowire types have a single upstream PSP flow. Each upstream PSP flow applies to a level of QoS in the CIN. This specification does not specify how the RPD schedules upstream traffic from multiple PSP flows or multiple sessions onto an upstream CIN.

6.1.1 DOCSIS D-MPT Data Path

The DEPI DOCSIS D-MPT flows contain DOCSIS frames using the format 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 RPD searches the D-MPT or MCM 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.

In D-MPT mode, MPEG packets are received by the RPD 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 PSP is a Layer 3 convergence layer protocol, which allows packets to be consecutively streamed together and fragmented at arbitrary boundaries. The primary rationale for PSP mode is to facilitate Quality of Service. This mode is to be used for transporting traditional DOCSIS data and signaling messages that 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; see Section 7.4.2.1.1 for more information.

The RPD MUST support reception and reassembly of PSP packets on at least four (4) downstream PSP flows per PSP session. The RPD MUST support transmission of at least four (4) upstream PSP flows for upstream data pseudowire types and one (1) PSP flow for all other pseudowire types.

6.1.2.1 PSP Downstream Data Path Packet Scheduler

Each downstream 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 CCAP Core and RPD) of the PSP flow that carried the DOCSIS frames. The Packet Scheduler is also responsible for inserting DOCSIS SYNC messages within the time interval specified by the GCP protocol. The Packet Scheduler SHOULD support a strict priority scheduler. The Packet Scheduler MAY support other queue scheduling logic.

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 is Weighted Fair Queuing (WFQ). Downstream packet scheduling in the RPD should not be confused with the more complex DOCSIS upstream scheduler, which deals with both requests and grants.

For SC-QAM channels, the output of the Packet Scheduler goes to a Transmission Convergence engine that places the DOCSIS frames into MPEG packets according to the requirements in [DRFI]. This includes the insertion of stuff bytes and the DOCSIS SYNC message as described by Section 6.1.3. For OFDM channels, the output of the Packet Scheduler is sent to the OFDM Transmission Convergence engine, which is explained in [PHYv3.1]. The output of the SC-QAM or OFDM Transmission Convergence (TC) engine is sent to the modulator and further to the DS RF Interface.

PSP mode provides acceleration of MAPs through the network, in an attempt to reduce Request Grant latency. The PSP mode is most useful when all or nearly all 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 technology.

For downstream PSP sessions carrying DOCSIS traffic, the CCAP Core SHOULD transmit MAPs and UCDs on the same downstream PSP flow that is assigned to the Expedited Forwarding (46) PHB-ID.

The RPD downstream packet scheduler MUST support at least two QoS queues on which it enqueues packets from PSP downstream flows assigned to the Expedited Forwarding (46) and Best Effort (0) PHB-IDs, respectively. The RPD SHOULD provide strict high priority scheduling to the QoS queue for PSP flows assigned to the Expedited Forwarding (46) PHB-ID. Alternatively, the RPD can signal the capability to support direct mapping between PSP Flows to strict priority QoS queues. When the RPD supports direct mapping of downstream flows to strict priority queues, the PSP flow with lowest numbered ID is mapped to the lowest priority QoS queue and the PSP flow with highest numbered ID is mapped to the highest priority QoS queue. When the RPD supports direct mapping of downstream flows to strict priority queues then the signaling of PHB-IDs is ignored.

6.1.3 DOCSIS SYNC Message

6.1.3.1 SYNC Message Format

The DOCSIS Time Synchronization (SYNC) MAC message is transmitted by a CCAP Core 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 5.




Figure 5 - Format of a DOCSIS SYNC MAC Message

The CCAP Core MUST send a MAC message with fields defined as shown in Table 2.

Table 2 - DOCSIS SYNC MAC Message Fields

Field	Description
FC, MAC PARM, LEN, HCS	Common MAC frame header with FC_PARM field to indicate a timing header; refer to [RF1v2.0] for details.
Destination Address (DA)	Always set to the DOCSIS MAC multicast address of 01-E0-2F-00-00-01.
Source Address (SA)	The MAC address of the CCAP Core. In PSP mode, the RPD learns the appropriate CCAP Core MAC address via explicit signaling during GCP 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 [ISO 8802-2].
SSAP	The LLC null source SAP (00), as defined by [ISO 8802-2].
Control	Unnumbered information frame (03), as defined by [ISO 8802-2].
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 RPD MUST derive a local DOCSIS timestamp as specified in [R-DTI]. In the D-MPT or MCM mode, the RPD MUST be capable of correcting all embedded SYNC messages in the DOCSIS stream. For the DOCSIS PSP mode, the RPD MUST be capable of inserting DOCSIS SYNC messages based on its internal timestamp into the downstream MPEG-TS stream. In PSP mode, the RPD MUST insert the SYNC message starting at the sixth byte of the MPEG-TS frame (the fifth byte is the MPEG pointer field). When referencing a timebase to use for SYNC timestamp insertion or correction, the RPD 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 at the R-DTI Local PTP Clock, as specified in [R-DTI].

This time difference between the time at the R-DTI Local PTP Clock and the time reference applied by an RPD is expected to be essentially constant. All timing and jitter requirements that are stated in this specification and [DRFI] still apply. The requirements here do not preclude an RPD from taking longer than 100 clock cycles to process the R-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 R-DTI conveyed timebase.

When using the D-MPT mode, the CCAP Core MUST generate SYNC messages and include them in the D-MPT payload for DOCSIS channels. The CCAP Core MUST insert the SYNC message starting at the sixth byte of the MPEG-TS frame (the fifth byte is the MPEG pointer field). This simplifies 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 CCAP Core is not required to accurately reflect the current timestamp received via an R-DTI Client in the CCAP Core. For example, it is permissible for the CCAP Core to use a value of zero for the CMTS timestamp in all SYNC messages. When using the DOCSIS PSP mode, the CCAP Core MUST NOT generate SYNC messages as part of the PSP payload.

6.1.4 Latency and Skew Requirements for DOCSIS Channels

6.1.4.1 Latency

For PSP DEPI DOCSIS 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 RPD DEPI port to the time that the first bit of the DOCSIS MAC frame exits the RPD RFI port. For D-MPT and MCM DEPI sessions, latency is defined as the absolute difference in time from when the last bit of a DEPI packet enters the RPD DEPI port to the time that the first bit of the first MPEG packet contained within said DEPI packet exits the RPD RFI port. At the RPD input, the last bit of the arriving DEPI packet is used because the RPD's Layer 2 interface (e.g., GigE or 10 GigE port) must receive the entire packet before the RPD can begin processing.

At the RPD output, the first bit of the first MPEG packet (in D-MPT or MCM 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 following paragraphs). 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.

In PSP DEPI sessions, the multiple flows supported in the RPD provide for prioritized access to the channel. This specification defines a concept of "isolated packet stream" for the purpose of establishing a requirement of forwarding latency through an RPD in the case when the internal buffering and queuing effects on latency are smallest. Isolated packets are spaced such that, at the nominal downstream data rate, the RPD would complete transmission of the current packet before the arrival of the next packet.

In the absence of higher priority traffic, and regardless of the amount of lower priority traffic, the RPD **MUST** forward isolated packets in each DEPI flow with a latency of less than 200 μ s plus the delay of the interleaver.

NOTE: The value stated in the above requirement accounts for the OFDM channel codeword builder delays.

The "isolated packet" forwarding delay requirements, outlined above, are applicable only to the forwarding of DOCSIS packets. These requirements are not applicable to forwarding of video, or any other type of traffic.

6.1.4.2 Skew

Skew is defined as the difference between the maximum latency and the minimum latency through the RPD, as measured from two reference bits on the network interface to the same bits on two separate RF outputs. Skew is measured with the PHY parameters set equally on the downstream channels under measurement.

The skew between any two bonded downstream channels within an RPD **MUST** be less than 500 μ s plus the delay introduced by the OFDM channel codeword builder, if applicable. The skew requirements for the RPD are implicitly met when the RPD meets the latency requirements described in Section 6.1.4.1. This requirement is intended for the transmission of skew-sensitive traffic such as bonded traffic.

6.1.4.3 Buffering Requirements for DOCSIS Channels

The RPD **MUST** allow enough buffer memory per DOCSIS channel to buffer at least 8 ms worth of traffic across all L2TPv3 sessions destined for that DOCSIS channel.

6.1.4.4 Buffering Requirements for Video Channels

For D-MPT video sessions, latency is defined as the absolute difference in time from when the last bit of a DEPI packet enters the RPD DEPI port to the time that the first bit of the first MPEG packet contained within said DEPI packet exits the RPD RFI port. The RPD **MUST** allow enough buffer memory per video QAM channel to buffer at least 10 msec worth of traffic for that video QAM channel.

6.2 System Timing Considerations

System timing considerations are described in [R-DTI].

6.3 MPEG-TS Requirements

The CCAP Core MUST support the [EQAM-VSI] specification. The output of the CCAP Core MUST be a constant bit rate MPTS with NULL packets inserted as needed for video channels and PCR timestamps that comply with MPEG specifications. Null packet insertion for DOCSIS channels is optional. The output MPTS is then transported using the D-MPT pseudowire.

The CCAP Core MUST support the D-MPT pseudowire type for MPEG-TS video packets.

The RPD MUST support the D-MPT pseudowire type for MPEG-TS video packets.

The CCAP Core MAY support the MCM pseudowire type for MPEG-TS video packets.

The RPD MAY support the MCM pseudowire type for MPEG-TS video packets.

The CCAP Core MUST be capable of being synchronized with the R-DTI clock.

The CCAP Core MUST generate an MPTS with the exact same data rate resulting from the M/N rate or symbol rate it configures at the RPD for the QAM channel to which the MPTS is destined.

The CCAP Core SHOULD comply with [RMI-ERM-EDGE].

The CCAP Core MUST participate in session signaling for setup and teardown of VoD and SDV sessions.

The RPD MUST support Synchronous mode.

The RPD SHOULD support Asynchronous mode.

If the RPD supports Asynchronous mode, then the RPD MUST have an MPEG-compliant clock per [ISO 13818-1]. Detailed clock requirements for the RPD can be found in [R-DTI].

If the RPD supports Asynchronous mode, then the RPD MUST insert or delete NULL packets as needed to create an MPEG-compliant output MPTS. PCR timestamps at the output of the RPD MUST be in compliance with [ISO 13818-1].

Although Synchronous mode is specified as a requirement, it is possible that operators may want to deploy systems without a 1588 source for an EQAM Core. In such cases it is assumed that operators are using RPDs that are capable of operating in Asynchronous mode.

The RPD advertises its capability to support Asynchronous mode as the number of capable QAM channels per RF port basis to the CCAP Core, via the GCP protocol. The CCAP Core provides an ability to configure the RPD mode of operation (Synchronous or Asynchronous) via GCP. When configuring the RPD for Asynchronous mode operation, the CCAP Core should take into account the capability advertised by the RPD and ensure that the RPD does not exceed its capability.

6.4 Multicast DEPI Requirements

CCAP Core SHOULD support PIM, IGMPv3, and MLDv2 on the DEPI interfaces. The CCAP Core SHOULD support the DEPI Remote Multicast Join AVP and the DEPI Remote Multicast Leave AVP.

The RPD MUST support IGMPv3 and MLDv2 on the CIN-facing interfaces. The RPD MUST support the DEPI Remote Multicast Join AVP and the DEPI Remote Multicast Leave AVP. On receiving a DEPI Remote Multicast Join AVP, the RPD MUST join the signaled multicast stream via IGMPv3/MLDv2. The RPD MUST be capable of joining 16 multicast streams simultaneously for each downstream RF port supported by the RPD. On receiving a DEPI Remote Multicast Leave AVP, the RPD MUST leave the multicast stream associated with the session via IGMPv3/MLDv2.

7 REMOTE PHY CONTROL PLANE PROTOCOL (NORMATIVE)

The R-PHY Control Plane protocol described in this section is responsible for managing all types of pseudowires and features including DEPI (DOCSIS and video), UEPI, SCTE 55-1 and SCTE 55-2 OOB as well as NDR/NDF.

The R-PHY protocol control plane signaling is conducted via GCP and L2TPv3 control protocol. The R-PHY L2TP control protocol signaling is based on the L2TPv3 control protocol introduced in [DEPI]. When compared to the original DEPI, the scope of control plane functionality conducted via L2TP control protocol has been reduced. All higher level control primitives, such as QAM channel configuration parameter management, have been removed from the R-PHY L2TP control protocol.

The L2TPv3 control protocol is used exclusively for managing the parameters of the R-PHY data plane connections, L2TPv3 pseudowires, including DEPI, UEPI, and OOB pseudowires. All other aspects of control plane operation, including RPD capability negotiation and RPD configuration management, are now conducted using GCP protocol.

This specification intentionally follows the provisions of [RFC 3931]. This section provides some examples of how L2TPv3 signaling is used, and includes the extensions and interpretations of the [RFC 3931] specification as it applies to R-PHY.

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

The examples presented in the following sections are intended to represent selected aspects of the R-PHY topology. The R-PHY architecture is not limited only to the topologies shown in the examples. Other topologies-or combinations of topologies presented in the examples-can be supported by an R-PHY system.

7.1 Remote PHY Topology

Figure 6 shows how the Remote PHY architecture maps to L2TP topology in a scenario where a single CCAP Core device is attached to a single RPD.




Figure 6 - End-to-end Remote PHY Topology

In L2TPv3 nomenclature, the CCAP Core and the RPD are known as an L2TP Access Concentrator (LAC). The CCAP Core and the RPD are considered to be peers and may also be referred to as L2TP nodes or as L2TP Control Connection Endpoints (LCCEs). For the purposes defined in this specification, each LCCE is identified on the network by a single IP address. The connections between two LCCEs are known as Pseudowires (PWs).

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

In the scenario presented above in Figure 6, a single L2TP connection is established between the pair of LCCEs as well as a single GCP connection.

First a GCP connection is established between the two LCCEs, through which the RPD advertises its capabilities to the CCAP Core and through which the CCAP Core configures the parameters of the peer RPD. A single GCP connection is utilized to configure all parameters, including those within the scope of DEPI, UEPI, OOB channels, or R-DTI.

Next, an L2TP control connection is established between the two LCCEs, and sessions are established. An L2TP session is established before L2TP begins to forward session frames for data transmission. Multiple sessions, including DEPI, UEPI, and OOB, are typically bound to a single L2TP control connection.




Figure 7 - Remote PHY Topology with Separate CMTS Core and EQAM Cores

Figure 7 depicts an arrangement where the CCAP Core consists of three individual core devices. The DOCSIS service is facilitated by a CMTS Core which terminates all DOCSIS DEPI and UEPI tunnels. Two separate EQAM Cores provide video service:

- EQAM Core 1 supplies VoD service transmitted over a dedicated set of video QAM channels in the RPD.
- EQAM Core 2 produces a set of broadcast video channels for transmission over another set of video QAM channels of the RPD.

In the Figure 7 example, the RPD maintains three separate connection sets, one set for each core device. Each of the core devices has its own GCP connection and L2TP control connection. Each core device is responsible for managing the parameters of its associated channels and for the administration of its corresponding L2TP pseudowires.

An important peer multi-connectivity aspect of R-PHY topology is presented in Figure 8.




Figure 8 - Remote PHY Topology with Multiple Connections

In the Figure 8 example, both the CCAP Core and the RPD include multiple network ports through which R-PHY network traffic is transmitted and received. In the figure, each of the ports is represented by an LAC/LCCE with an individual IP address. As shown in Figure 9, two DEPI tunnels are established between different pairs of network ports of each device. In such a topology, a single GCP and two L2TP control connections are created to manage R-PHY parameters and L2TP data pseudowires.

RPDs support the optional ability to provide daisy-chain connectivity to other RPDs. A daisy chain is a connection of RPDs in a series, one after another. To support daisy chaining, the R-PHY has to provide at least one network port to which another RPD can be connected. Daisy chaining introduces another variation of the R-PHY topology.

7.2 Remote PHY Addressing

The set of network connections between a pair of devices consisting of one CCAP Core and one RPD includes

- Exactly one GCP connection per pair of CCAP Core and RPD, and;
- A number of L2TPv3 Control Connections: one L2TPv3 Control Connection per pair of interconnected L2TP Control Connection Endpoints (LCCEs), and;
- A number of L2TPv3 pseudowires (sessions), one or more per channel, and;
- Sessions utilizing the PSP pseudowire format may consist of multiple data flows.

An example of hierarchy in R-PHY connectivity is presented in Figure 9. In the figure, the CCAP Core contains two LCCEs (LCCE M and LCCE N) and the RPD contains 2 LCCEs (LCCE P and LCCE Q). Each LCCE represents two DEPI channels, an SC-QAM channel and an OFDM channel.

Figure 9 shows a single GCP connection between the two devices, two L2TP Control Connections, one L2TP Control Connection per pair of interconnected LCCEs, a DEPI data session for each channel, and two data flows for each data session.




Figure 9 - R-PHY Connectivity Hierarchy

The RPD MUST use the IP address of the CCAP Core to establish the GCP control connection during the initial system configuration. The RPD MUST establish exactly one GCP connection for each interconnected CCAP Core device.

The CCAP Core SHOULD use the IP address of the RPD and RPD Channel Selector to uniquely identify a QAM Channel within an RPD during L2TP session configuration. RPD Channel Selector is defined in Section 7.5.1.12

The CCAP Core MUST establish at most one L2TP control connection for each interconnected LCCE pair. This control connection manages all L2TP data sessions between the pair of LCCEs in the CCAP Core and the RPD. Since each device may be logically represented as multiple LCCEs, more than one L2TP connection can be created for a set comprising one CCAP Core and one RPD.

The session scaling requirements for UEPI and OOB are listed in [R-UEPI] and [R-OOB], respectively.

Each R-PHY session uses one of the Pseudowire types described in Table 8. Per [RFC 3931], both the CCAP Core and the RPD assign a unique L2TPv3 session ID for each session.

The CCAP Core MUST NOT attempt to create a session for a channel for which the CCAP Core already has an active session of the same type. Unless specifically configured otherwise, the RPD MUST reject an attempt to establish a session to a channel for which a session of same type has already been established. The requirements listed above are modified when RPD and the CCAP Core operate in redundant mode. For more information on R-PHY redundancy support, see DEPI Path Redundancy (Normative).

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

7.3 CCAP Core L2TP Control Message Format

7.3.1 L2TP Control Message without UDP Header

The properties of the L2TP control message header are shown in Figure 10.




Figure 10 - L2TP Control Message without UDP Header

7.3.2 L2TP Control Message with a UDP Header

The structure of the L2TP control message header is shown in Figure 11.




Figure 11 - L2TPv3 Control Packet with UDP Header

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. The operation of R-PHY protocols 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 CCAP Core, the Ethernet Destination Address will be the Ethernet address of the RPD or of the next hop router. When the RPD receives this frame, the Ethernet Source Address will be the Ethernet address of the output port of the CCAP Core or of the previous hop router.

If the networking interface is Ethernet, the CCAP Core MUST support the Ethernet header. If the networking interface is Ethernet, the RPD MUST support the Ethernet header. If another physical layer interface is used instead of Ethernet, then the Ethernet header 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]. This header provides frame prioritization and VLAN support at Layer 2.

The CCAP Core MUST be able to support the Ethernet 802.1Q header. The CCAP Core SHOULD support mapping of 802.1Q user priority of tunneled packets based on the DSCP value of a tunnel's IP packet. The RPD SHOULD support the Ethernet 802.1Q header.

7.3.3.3 IPv4 Header

The format of an IPv4 header is defined by [RFC 791].

Because of implementation considerations and for coexistence with network policies that are not amenable to IP fragmentation, the receiving devices (R-PHY entities in DEPI, CCAP Core in UEPI) are not required to perform IP reassembly. The CCAP Core MUST NOT use IP fragmentation. The CCAP Core MUST set the IP DF (Don't Fragment) bit in transmitted packets. The RPD MUST NOT use IP fragmentation. The RPD MUST set the IP DF (Don't Fragment) bit in transmitted packets.

The CCAP Core MUST support a configurable 6 bit Differentiated Services Code Point (DSCP) as defined by [RFC 2983] and [RFC 3260].

The CCAP Core MUST support the IPv4 header. The RPD MUST support the IPv4 header.

7.3.3.3.1 LCCE IP Addresses

A CCAP Core MAY initiate L2TPv3 from an LCCE entity IP address that differs from the GCP entity IP address with which it communicates to the RPD. A CCAP Core MAY initiate an L2TPv3 session to any IP address reported by the RPD as representing an LCCE entity, even if the address does not match the IP address with which the RPD maintains a GCP session with the CCAP Core.

The RPD MUST support CCAP Core LCCE entity IP addresses that differ from the CCAP Core GCP Entity. The RPD MUST support at least four (4) different values of the GCP LCCE entity IP address.

If IKEv2 mutual authentication was established for the CCAP Core and the RPD GCP entity IP addresses, the CCAP Core and RPD MUST establish derived security associations for each pair of LCCE entity IP addresses before attempting to start a control session between those LCCE IP addresses.

Both the CCAP Core and the RPD MUST use the same LCCE IP address for the control session and all data sessions started by that control session.

7.3.3.3.2 DSCP

The R-PHY OSSI specification defines a minimum standard policy by which the CCAP Core selects DSCP values for L2TPv3 packets sent by itself and by the RPD. The specification also permits the CCAP Core to have a vendor-specific L2TPv3 DHCP policy.

For L2TPv3 control packets, the CCAP Core MUST transmit with the DSCP specified by its L2TPv3 DSCP policy. When its L2TPv3 control DSCP is other than BestEffort(0), the CCAP Core MUST communicate that DSCP value to the RPD with a Control Connection DSCP AVP (see Section 7.5.1.6). When it receives a Control Connection DSCP AVP, the RPD MUST transmit all upstream L2TPv3 control packets with that DSCP. When the RPD does not receive a Control Connection DSCP AVP, the RPD MUST transmit upstream control packets with a DSCP of BestEffort(0).

The CCAP Core MUST transmit downstream L2TPv3 data sublayer packets with a DSCP selected by its L2TPv3 policy, e.g., as configured for a pseudowire type and PSP Flow. Note that this DSCP may or may not match the PHB-ID signaled to the RPD for the flow in the DEPI Resource Allocation Request AVP (Section 7.5.3.2). This applies for PSP and non-PSP pseudowires with downstream sublayer data (see Table 10 and Table 24).

For downstream L2TPv3 DLM sublayer packets (which include a downstream flow ID), the CCAP Core MUST transmit with the same DSCP value selected by its policy for the identified downstream flow. For L2TPv3 sessions with downstream flows (PSP or non-PSP) that are eligible for a DLM request, the CCAP Core SHOULD signal PHB-ID in a DEPI DLM Response DSCP AVP. The RPD MUST transmit DLM responses with a DSCP equal to the PHB-ID signaled in the DEPI DLM Response DSCP AVP for the session. If the CCAP Core does not signal PHB-ID in a DEPI DLM Response DSCP AVP, the RPD SHOULD transmit DLM Responses with the same DSCP as the corresponding requests.

S-BFD packets apply to a pseudowire as a whole and do not signal a flow ID in the L2TPv3 sublayer header. The CCAP Core MUST transmit S-BFD requests with a DSCP specified by its local DSCP policy.

For pseudowires eligible for S-BFD requests, the CCAP Core SHOULD signal PHB-ID in a DEPI S-BFD Return Path DSCP AVP. The RPD MUST transmit an upstream S-BFD response with a DSCP equal to the PHB-ID signaled in the DEPI S-BFD Return Path DSCP AVP for the session. If the CCAP Core does not signal PHB-ID in a DEPI S-BFD Return Path DSCP AVP, the RPD SHOULD transmit S-BFD Responses with the same DSCP as corresponding requests.

For pseudowires with upstream sublayer data (see Table 10 and Table 24), the CCAP Core MUST include the Upstream Flow AVP of the ICRQ control message it sends to establish the L2TPv3 data session with a PHB-ID equal to the DSCP set by its policy.

When an RPD receives an Upstream Flow AVP for a session, it MUST transmit upstream L2TPv3 sublayer data packets with a DSCP that is equal to the PHB-ID assigned in that AVP for each upstream flow ID.

7.3.3.4 IPv6 Header

The IPv6 header is defined by [RFC 2460].

The CCAP Core MUST support a configurable 6 bit Differentiated Service Code Point (DSCP) of the IPv6 Traffic Class as defined by [RFC 2983] and [RFC 3260].

The CCAP Core MUST support the IPv6 header. The RPD MUST support the IPv6 header.

The CCAP Core MUST support for IPv6 headers all requirements of Sections 7.3.3.3.1 and 7.3.3.3.2.

The RPD MUST support for IPv6 headers all requirements of Sections 7.3.3.3.1 and 7.3.3.3.2.

7.3.3.5 UDP Header

To enhance dispersion of traffic across Link Aggregation Groups (LAGs) in the CIN, R-DEPI supports the optional inclusion of a UDP header in L2TPv3 control and data traffic. Currently, most commercial CIN routers do not include the L2TPv3 session ID in the hash calculation that selects a particular LAG link, but almost all include the UDP source and destination port number in the hash calculation.

The CCAP Core MUST support L2TPv3 without a UDP header.

The CCAP Core MUST support L2TPv3 with a UDP header.

The RPD MUST support L2TPv3 without a UDP header.

The RPD MAY support L2TPv3 with a UDP header.

An RPD that supports UDP MUST indicate its capability to do so via GCP. A CCAP Core MUST use a UDP header for all L2TPv3 control and data packets when configured to use UDP headers via OSSI configuration, and all RPDs terminating a pseudowire report when they are capable. An RPD MUST use a UDP header for all L2TPv3 control and data packets when it is capable of supporting UDP headers and receives an initial SCCRQ control message from a CCAP that includes a UDP header.

7.3.3.5.1 Control Packet UDP Port Numbers

When using UDP headers, CCAP Core and RPD MUST negotiate a pair of ephemeral UDP ports for the LCCE control connection as follows:

1. The CCAP Core MUST send the initial SCCRQ to an RPD with the well-known destination UDP port 1701 (as assigned to L2TPv3 per [IANA-PORTS]) and an ephemeral source UDP port number denoted here as *core-control-port*;
2. The RPD MUST send its SCCRP reply with the destination UDP port set to *core-control-port* and an ephemeral source UDP port number denoted here as *rdp-control-port*;
3. The CCAP Core MUST send all control packets of an LCCE control connection after SCCRQ to the destination UDP port *rdp-control-port* with the source UDP port *core-control-port*;
4. The RPD MUST send all control packets of an LCCE control connection after SCCRP to the destination UDP port *core-control-port* and the source UDP port *rdp-control-port*.

The negotiation scheme outlined above is also described in [RFC 3931].

The CCAP Core SHOULD select UDP port 1701 for *core-control-port*.

The RPD SHOULD select UDP port 1701 for *rdp-control-port*.

7.3.3.5.2 Downstream Data UDP Port Numbers for Downstream Sessions

When using UDP headers, the CCAP Core and RPD MUST negotiate the destination UDP port number for each downstream L2TPv3 data session with the following requirements:

For unicast downstream data sessions, the RPD MAY select a UDP port number of its own choosing and communicate it to the CCAP Core via DEPI Local UDP Port AVP in ICRP message. The CCAP Core MUST use the UDP port number selected by the RPD as the destination UDP port number in all packets sent to the RPD on the session. If RPD does not communicate its selection of the UDP port in ICRP, the CCAP Core MUST use *rdp-control-port* as the destination UDP port number in all packets sent to the RPD on the session.

The method for allocation of destination UDP ports for multicast sessions differs from the unicast method. The protocol allows the transmitter, i.e., the CCAP Core, to select UDP port for all recipients, i.e., RPDs, of the multicast session. To permit a range of implementations and to avoid allocation conflicts between multiple CCAP Cores connected to same RPD, the CCAP Core supports operator configuration of Multicast UDP Port Range. The CCAP Core communicates Multicast UDP Port Range to the RPD via GCP. For multicast downstream data session, the CCAP Core MUST select a UDP port number from the Multicast UDP Port Range and communicates this number via DEPI Multicast UDP Port AVP to all RPDs in the multicast group. The RPD MUST NOT allocate UDP port numbers from the Multicast UDP Port Range for unicast communication.

For both unicast and multicast downstream data sessions the CCAP Core MAY select a source UDP port number of its own choosing for downstream L2TPv3 data packets and communicate it to the RPD via DEPI Local UDP Port AVP in ICRQ message. The CCAP Core MUST use the same source UDP port for all downstream data packets of a particular session. When the CCAP Core communicates the source UDP port number to the RPD via DEPI Local UDP Port AVP, the RPD MUST use this port as a destination UDP port number in all packets sent to the CCAP Core on the session (such as DLM and BFD). If the CCAP Core does not communicate its selection of the UDP port in ICRQ, the RPD MUST use *core-control-port* as the destination UDP port number in all packets sent to the CCAP Core on that session.

7.3.3.5.3 UDP Port Numbers for Upstream Sessions

When using UDP headers, the CCAP Core and RPD negotiate the destination UDP port number for each upstream session as follows:

For upstream data sessions the CCAP Core MAY select UDP port number of its own choosing and communicate it to the RPD via DEPI Local UDP Port AVP in ICRQ message. The RPD MUST use the UDP port number selected by the CCAP Core as the destination UDP port number in all packets sent to the CCAP Core on the session. If

CCAP Core does not communicate its selection of the UDP port in ICRQ, the RPD MUST use *core-control-port* as the destination UDP port number in all packets sent to the CCAP Core on the session.

For upstream data sessions the RPD MAY select a source UDP port number of its own choosing for L2TPv3 data packets and communicate it to the CCAP Core via DEPI Local UDP Port AVP in ICRP message. The RPD MUST use the same source UDP port for all upstream data packets of a particular session. When the RPD communicates the source UDP port number to the CCAP Core via DEPI Local UDP Port AVP, the CCAP Core MUST use this port as a destination UDP port number in all packets sent to the RPD on the session (such as BFD). If the RPD does not communicate its selection of the source UDP port in ICRP, the CCAP Core MUST use *rpd-control-port* as the destination UDP port number in all packets sent to the RPD on that session.

7.3.3.5.4 UDP Checksum

When L2TPv3 Control Message Authentication is not in use and UDP headers are enabled, both the RPD and CCAP Core MUST support UDP checksum verification as defined in [RFC 768] for L2TPv3 control packets. The sender of a control packet MUST calculate and populate the UDP checksum as per [RFC 768]. The receiver of a control packet MUST support validation of a nonzero UDP as per [RFC 768]. The receiver of a control packet MUST reject a packet with a 0 checksum or one that fails UDP checksum verification.

When L2TPv3 Control Message Authentication is in use, the UDP checksum verification is redundant. Therefore, when L2TPv3 Control Message Authentication is in use and UDP headers are enabled, the sender of a control packet MAY set the UDP checksum field to 0, and the receiver MUST ignore a UDP checksum of 0.

For L2TPv3 data packets, including DLM and BFD, the sender (CCAP Core or RPD) SHOULD set the UDP checksum to 0. For L2TPv3 data packets, the receiver (CCAP Core or RPD) MUST ignore a UDP checksum of 0 and MAY ignore a UDP checksum that is nonzero. **UDP Checksum and IPv6 transport**

In general, [RFC 2460] mandates that IPv6 nodes operating with UDP datagrams include a valid UDP checksum in transmitted datagrams, and discard received IPv6 datagrams in which UDP checksum is NULL. [RFC 6935] relaxes these requirements for tunneled data protocols and permits a zero UDP checksum in the outer encapsulation of IPv6 datagrams of a tunnel protocol. This relaxation does not apply to L2TPv3 control messages. The sender of L2TPv3 control messages with IPv6 transport is required to calculate and insert the adequate UDP checksum into the UDP header.

The CCAP Core MUST calculate the UDP checksum and insert the checksum into the UDP header of all transmitted L2TPv3 control messages with IPv6 transport and UDP header.

The RPD MUST calculate the UDP checksum and insert checksum into the UDP header of all transmitted L2TPv3 control messages with IPv6 transport and UDP header.

The CCAP Core MUST verify the UDP checksum in received L2TPv3 Control Messages with non-zero UDP checksum.

The RPD MUST verify the UDP checksum in received L2TPv3 Control Messages with non-zero UDP checksum.

The CCAP Core MAY calculate the UDP checksum and insert the checksum into the UDP header in transmitted L2TPv3 data messages with IPv6 transport and UDP header.

The RPD MAY calculate the UDP checksum and insert the checksum into the UDP header in transmitted L2TPv3 data messages with IPv6 transport and UDP header.

The CCAP Core SHOULD verify the UDP checksum in received L2TPv3 Data Messages with non-zero UDP checksum.

The RPD SHOULD verify the UDP checksum in received L2TPv3 Data Messages with non-zero UDP checksum.

7.3.3.6 CRC

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

The CCAP Core MUST support the CRC field. The RPD 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. Table 3 describes the significance of the fields.

Table 3 - L2TPv3 Control Header Fields

Field	Description
T	Type bit. (See requirement 1 immediately following this table)
L	Length bit (See requirement 2 below)
S	Sequence bit. (See requirement 3 below)
X	Reserved bits. (See requirement 4 below)
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 (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.

Requirements referenced in Table 3 above:

1. The L2TPv3 Control Header Field Type T bit MUST be set to 1, indicating that this is a control message.
2. The L2TPv3 Control Header Field Length L bit MUST be set to 1, indicating that the Length field is present.
3. The L2TPv3 Control Header Field Sequence S bit MUST be set to 1 indicating that sequence numbers (Ns and Nr) are present.
4. All reserved L2TPv3 Control Header Field bits MUST be set to 0 on outgoing messages and be ignored on incoming messages.

7.3.4.2 Attribute Value Pairs (AVP)

There can be one or more attribute value pairs (AVPs) per R-PHY control message. Table 4 describes the significance of the fields.

Table 4 - R-PHY Control Message AVPs

Field	Description
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 3931].
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 zeroes 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 L2TP Signaling

The supported L2TPv3 messages for the R-PHY control plane are shown in Table 5.

Table 5 - R-PHY L2TP Control Messages

Mnemonic		Name
Control Connection Management		
1	SCCRQ	Start-Control-Connection-Request
2	SCCRP	Start-Control-Connection-Reply
3	S CCCN	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.

A reliable control message delivery mechanism 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 specified by Control Message Timeout (refer to Parameters and Constants), then the control message MUST be retransmitted up to the value of Control Message Retry Count (refer to Parameters and Constants). 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 each in this scheme.

Control message authentication MAY be supported. If control message authentication is supported, the methods described in section 5.4.1 of [RFC 3931] SHOULD be followed.

The following sections show flow diagrams that 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 12 shows the DEPI message exchange that sets up the control connection.




Figure 12 - R-PHY L2TP 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. Note that any L2TP Control Connection setup follows establishment of the GCP connection, including capability identification and device configuration. Each control connection has a Control Connection ID that is assigned by the recipient, and negotiated with Connection ID AVPs during the creation of a control connection.

The CCAP Core MUST support the ability to initiate control connection signaling (L2TPv3 caller). The RPD MUST support the ability to receive incoming control connection requests from the CCAP Core (L2TPv3 callee).

Establishment of L2TP control connections by the RPD is not required for the R-PHY protocol and is outside the scope of this specification.

7.4.1.2 Control Connection Teardown

Figure 13 shows the DEPI message exchange that tears down the control connection.




Figure 13 - R-PHY L2TP 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 message. Thus, it is not necessary to clear each session individually when tearing down the entire control connection. The peer receiving the StopCCN message **MUST** maintain session and control state for a period of time equal to StopCCN Timeout (see Parameters and Constants) 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 receives a 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 that sent the StopCCN message.

7.4.1.3 Control Connection Keep-Alive

Figure 14 shows the DEPI message exchange that refreshes the control connection.




Figure 14 - R-PHY L2TP 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 Figure 14) has passed without receiving any message (data or control) from the peer.

7.4.2 Session Signaling

7.4.2.1 Session Setup

Figure 15 shows the DEPI message exchange that sets up a DOCSIS session.




Figure 15 - R-PHY L2TP Session Setup

After successful control connection establishment, individual sessions may be created. Each session corresponds to a single data stream between the two LCCEs. The R-PHY protocol introduces new pseudowire types, some of which tunnel data from the RPD to the CCAP Core. In addition to the mandatory and optional AVPs in [RFC 3931], the DEPI-specific AVPs shown in Figure 15 are used as part of the session setup.

The ICRQ message contains the Remote End ID AVP that contain one or more RPD Channel Selectors of the Channel(s) for which the session is intended. Multiple RPD Channel Selectors can be included for those sessions that carry data for multiple channels. RPD Channel Selector is defined in Section 7.5.1.12.

The Session ID selection method and protocol signaling differs between unicast and multicast DEPI sessions.

For unicast DEPI sessions, the CCAP Core sets the Remote Session ID AVP in the ICRQ message to zero per [RFC 3931].

For multicast DEPI sessions, the CCAP Core allocates the session ID from the multicast session ID pool and sends its value in the Remote Session ID AVP in an ICRQ message. The multicast session ID pool consists of a set of reserved ID values in a range from 0x80000001 to 0x8000FFFF. The range of the multicast session ID pool can be further restricted by the operator via administrative controls to prevent two CCAP Cores serving overlapping sets of RPDs from using the same set of multicast session IDs.

The CCAP Core **MUST** ensure that it assigns a set of unique multicast Session ID values to any of its attached RPDs.

The CCAP Core **MUST NOT** select session ID values from the multicast session ID pool range for unicast DEPI sessions. The RPD **MUST NOT** select session ID values from the multicast session ID pool range for unicast DEPI sessions.

These requirements are intended to allow for RPD implementations in which individual L2TPv3 sessions can be identified by the Session ID alone, without relying on protocol identifiers from other layers.

7.4.2.1.1 ICRP, ICRQ, and ICCN Messages

The ICRP message contains the Local Session AVP that indicates the session ID the RPD wants to use. As mentioned earlier, the RPD does not allocate its session ID for multicast DEPI sessions. In an ICRP message related to a multicast DEPI session, the value of Local Session ID in that ICRP message is the same value which was provided by the CCAP Core to the RPD via Remote Session ID AVP in the ICRQ message. Note that the ICRP does not contain a suite of QAM Channel AVPs that indicate the RPD's current configuration and whose parameters can be changed. These parameters are negotiated prior to L2TP session establishment via GCP. The ICRQ message for R-PHY protocol can still include the DEPI Resource Allocation AVP needed for management of DSCP values for the pseudowire.

The receipt and processing of the ICCN triggers the initiation of data forwarding within the RPD for the channel associated with the session. The RPD **SHOULD NOT** buffer data before the channel corresponding to the activated session is ready for data forwarding.

The RPD **SHOULD** set the Circuit Status AVP in the ICRP message to the down state. After the ICCN is received and when the channel associated with the session is operational, then the RPD **SHOULD** send an SLI message to the CCAP Core with the Circuit Status AVP set to the up state (refer to Section 7.5.1.5).

The CCAP 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 CCAP Core **MUST** be able to generate session establishment signaling. The RPD **MUST** support the ability to receive incoming session establishment requests from the CCAP Core.

NOTE: Establishment of L2TP sessions by the RPD is outside the scope of this specification.

7.4.2.2 Session Teardown

Figure 16 shows the DEPI message exchange that tears down a DOCSIS session.




Figure 16 - R-PHY 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 a StopCCN message. Thus, it is not necessary to clear each session individually when tearing down the entire control connection.

7.4.2.3 Session Updates

Figure 17 shows the DEPI message exchange that updates a DOCSIS session.




Figure 17 - R-PHY Session Updates

If there is a configuration change to one of the pseudowire parameters at the CCAP Core that is described by a DEPI AVP, the CCAP Core MUST send the updated AVP to the RPD with the Set-Link-Info (SLI) message. If there is a configuration change to one of the RPD parameters that is described by a DEPI AVP, the RPD MUST send the updated AVP to the CCAP Core using an SLI message.

7.4.3 Mandatory and Optional AVPs

In addition to the mandatory and optional AVPs listed in [RFC 3931] and modified by Table 7, the DEPI AVPs listed in Table 6 MUST be present in the R-PHY L2TP control message if mandatory, and MAY be present in R-PHY L2TP control messages if optional.

Table 6 - R-PHY L2TP Mandatory and Optional AVPs[16]

R-PHY L2TP Control Message	DEPI Mandatory AVPs	DEPI Optional AVPs
ICRQ	DEPI Resource Allocation Request(1) DEPI Local MTU Upstream Flow(2)	DEPI Local UDP Port DEPI Multicast UDP Port DEPI DLM Response DSCP DEPI S-BFD Return Path DSCP
ICRP	DEPI Resource Allocation Reply(1) DEPI Remote MTU	DEPI Local UDP Port
ICCN		
SCCRQ	DEPI Multicast Capability	
SCCRP	DEPI Multicast Capability	
CDN		DEPI Result Code
StopCCN		DEPI Result Code

R-PHY L2TP Control Message	DEPI Mandatory AVPs	DEPI Optional AVPs
SLI		DEPI Result Code
Note 1. Required for downstream sublayer sessions only. Note 2. Required for upstream sublayer sessions only.		

7.5 AVP Definition Updates

7.5.1 Conventional L2TPv3 AVPs

The AVP types from [RFC 3931] and [RFC 3308] that are supported as part of this specification are listed in Table 7.

Table 7 - R-PHY Supported L2TPv3 AVPs

Attribute Type	Control, Session	Description	Required	Not Required
0	C, S	Message Type	♦	
1	S	Result Code	♦	
5	C, S	Control/Session Tie Breaker		♦
7	C	Host Name	♦	
8	C	Vendor Name		♦
10	C	Receive Window Size		♦
15	S	Serial Number	♦	
25	S	Physical Channel ID		♦
34	S	Circuit Errors		♦
36	C, S	Random Vector		♦
47	C	Control Connection DSCP	♦	
48	S	Session DSCP		♦
58	C, S	Extended Vendor ID AVP		♦
59	C, S	Message Digest		♦
60	C	Router ID	♦	
61	C	Assigned Control Connection ID	♦	
62	C	Pseudowire Capabilities List	♦	
63	S	Local Session ID	♦	
64	S	Remote Session ID	♦	
65	S	Assigned Cookie		♦
66	S	Remote End ID	♦	
68	S	Pseudowire Type	♦	
69	S	L2-Specific Sublayer	♦	
70	S	Data Sequencing	♦	
71	S	Circuit Status	♦	
72	C	Preferred Language		♦
73	C	Control Message Authentication Nonce		♦
74	S	Tx Connect Speed		♦
75	S	Rx Connect Speed		♦

The CCAP Core MUST support the required L2TPv3 AVPs as designated in Table 7 above.

The RPD MUST support the required L2TPv3 AVPs as designated in Table 7 above.

Conventional AVPs whose usage is specific to the R-PHY protocol are described below. A more complete description, as well as requirements for the conventional AVPs, can be found in [RFC 3931].

The CCAP Core and RPD MUST transmit AVP fields marked as "reserved" with zero (0). The CCAP Core and RPD MUST ignore AVP fields marked as "reserved".

7.5.1.1 Message Type (All Messages)

Figure 18 shows the Message Type AVP.




Figure 18 - Message Type AVP

This AVP identifies the particular type of L2TPv3 Control Message that is being sent. It is always the first AVP.

7.5.1.2 Result Code (StopCCN, CDN)

Figure 19 shows the Result Code AVP.




Figure 19 - Result Code AVP

This message contains results codes, optional error codes, and optional error messages when tearing down a control connection or a session.

7.5.1.3 Host Name (SCCRQ, SCCRQ)

Figure 20 shows the Host Name AVP.




Figure 20 - Host Name AVP

The host name is typically the fully qualified domain name (FQDN) of each device.

7.5.1.4 Vendor Name (SCCRQ, SCCRQ)

Figure 21 shows the Vendor Name AVP.

7.5.1.8 Assigned Control Connection ID (SCCRQ, SCCRP, StopCCN)

Figure 25 shows the Assigned Control Connection ID AVP.




Figure 25 - Assigned Control Connection ID AVP

This is the assigned control connection ID for the R-PHY protocol. During SCCRQ, the CCAP Core uses this AVP to inform the RPD what value of Control Connection ID to use in the L2TPv3 Control Header for control messages originated by the RPD. During SCCRP, the RPD uses this AVP to inform the CCAP Core what value of Control Connection ID to use in the L2TPv3 Control Header for control messages originated by the CCAP Core. Since the CCAP Core has not been informed by the RPD prior to the first SCCRQ message, the CCAP Core uses a value of 0 for the control connection ID in the L2TPv3 Control Header of the SCCRQ message.

7.5.1.9 Pseudowire Capabilities List (SCCRQ, SCCRP)

Figure 26 shows the Pseudowire Capabilities List AVP.



Figure 26 - Pseudowire Capabilities List AVP

The Pseudowire Capabilities List AVP indicates the capabilities of the CCAP Core and RPD. There are two main PW types defined for an RPD and a CCAP Core; refer to Table 8 for details. DEPI also defines a set of pseudowire subtypes, which are communicated through a DEPI specific DEPI Pseudowire Subtype Capabilities List AVP.

Table 8 - Pseudowire Types

Pseudowire (PW) Type	Mnemonic	Value
MPT Pseudowire	MPTPW	0x000C*
PSP Pseudowire	PSPPW	0x000D*

The CCAP Core MUST indicate its support for pseudowire types by including one or more of the PW Types in the Pseudowire Capabilities List. The RPD MUST indicate its support for pseudowire types by including one or more of the PW Types in the Pseudowire Capabilities List.

7.5.1.10 Local Session ID (ICRQ, ICRP, ICCN, CDN, SLI)

Figure 27 shows the Local Session ID AVP.




Figure 27 - Local Session ID AVP

When a session is established, the CCAP Core and RPD each choose their own session ID and advertise it to each other with this AVP. This means that a single session setup sets up two unidirectional sessions, one in each direction.

7.5.1.11 Remote Session ID (ICRQ, ICRP, ICCN, CDN, SLI)

Figure 28 shows the Remote Session ID AVP.




Figure 28 - Remote Session ID AVP

When the CCAP Core or RPD send a session message to each other, the Remote Session ID is set to the session ID learned previously from the Local Session ID AVP. If the Remote Session ID is not yet known, it is set to 0.

The CCAP Core allocates session ID values for multicast DEPI sessions from the multicast session ID pool. The CCAP Core MUST send the Remote Session ID value in the ICRQ message when signaling multicast DEPI sessions.

7.5.1.12 Remote End ID (ICRQ)

Figure 29 shows the Remote End ID AVP.




Figure 29 - Remote End ID AVP

The Remote End ID AVP in R-DEPI has a value field which consists of a 2 octet long reserved field and one or more of 4 octet long array entries. With N array entries, the total length of the Remote End AVP can be expressed as 8+N*4.

Each array entry is a combination of two fields. These fields are RPD Channel Selector and a Channel ID from the PSP segment header (or MPTS_TAG from the DEPI-MCM header).

The RPD Channel Selector is a three-octet field. The RPD Channel Selector uniquely identifies a channel in the RPD. The inclusion of the Channel Selector is necessary to bind an L2TPv3 session to a channel or a sub-channel in the RPD. The structure of the RPD Channel Selector is presented on Figure 30.



Figure 30 - RPD Channel Selector

The fields of RPD Channel Selector are described in Table 9.

Table 9 - RPD Channel Selector

Field	Description
RF Port Index	8 bits. The index of the RPD's RF Port to which a channel or a sub-channel belongs. This field uniquely identifies US or Ds RF port in the RPD. The selection of US or DS RF port depends on Channel Type field.
Channel Type	8 bits. 0 - reserved 1 - DS-OFDM 2 - DS-OFDM-PLC 3 - DS-SCQAM 4 - US-ATDMA 5 - US-OFDMA 6 - SCTE-55-1-FWD 7 - SCTE-55-1-RET 8 - SCTE-55-2-FWD 9 - SCTE-55-2-RET 10 - NDF 11 - NDR
Channel Index	8 bits. An index - an identifier of a channel or a sub-channel of the selected type.

The last octet of each array entry in Remote End Id AVP value field is the Channel Id field. For pseudowires carrying data for multiple channels this field associates the channel indicated by the RPD Channel Selector with the Channel Id carried in the PSP segment table shown in Figure 56 or with the MPTS Tag in DEPI-MCM Sublayer header shown in Figure 54. The Channel Id field is not used when signaling all other DEPI pseudowire types. In these cases The Channel Id is transmitted as zero and ignored upon reception.

The CCAP Core MAY include the Remote End Id AVP multiple times in an ICRQ message when the CCAP Core signals a request to create a pseudowire carrying data for multiple channels. In such cases, the CCAP Core MUST include the Remote End Id AVP exactly once for each included channel or PLC-sub-channel. In all other cases, the CCAP Core MUST include a Remote End ID AVP exactly once in the ICRQ message.

7.5.1.13 Pseudowire Type (ICRQ)

Figure 31 shows the Pseudowire Type AVP.




Figure 31 - Pseudowire Type AVP

R-DEPI uses the Pseudowire Type values defined in Table 8 to indicate the major type of DEPI session requested. Additional information about the subtype of the DEPI Session requested is conveyed via DEPI Pseudowire Subtype AVP defined in Section 7.5.3.11.

7.5.1.14 L2-Specific Sublayer (ICRQ, ICRP, ICCN)

Figure 32 shows the L2-Specific Sublayer AVP.




Figure 32 - L2-Specific Sublayer AVP

The CCAP Core MUST include the L2-Specific Sublayer AVP in the ICRQ, ICRP and ICCN messages, indicating the L2-Specific Sublayer Header Type consistent with the Pseudowire type for the DEPI flow. The RPD MUST include the L2-Specific Sublayer AVP in the ICRP message, indicating the L2-Specific Sublayer Header Type consistent with the Pseudowire type for the R-PHY protocol flow. Additional information about the L2 Specific Sublayer is conveyed via DEPI L2-Specific Sublayer Subtype AVP described in Section 7.5.3.12.

Although L2TPv3 defines sessions as bidirectional, the Remote PHY architecture defines its sublayers with a single direction of downstream or upstream data flow, as indicated in Table 10 and further in Table 24. Note that DLM and S-BFD packets are considered (and marked in the sublayer header) as sublayer control packets, not sublayer data packets, and are transmitted in both directions.

Table 10 - L2-Specific Sublayer Types

L2-Specific Sublayer Type	Value	Direction
MPT Specific Sublayer	3*	Downstream or Upstream
PSP Specific Sublayer	4*	Downstream or Upstream

7.5.1.15 Data Sequencing (ICRP)

Figure 33 shows the Data Sequencing AVP.




Figure 33 - Data Sequencing AVP

The RPD MUST include the Data Sequencing AVP in the ICRP message, indicating Data Sequencing Level 2 is required (all incoming data packets require sequencing).

7.5.1.16 Circuit Status (ICRQ, ICRP, ICCN, SLI)

Figure 34 shows the Circuit Status AVP.




Figure 34 - Circuit Status AVP

Table 11 - Circuit Status AVP Details

Field	Description
N	1 bit. New bit. Indicates whether the status indication is for a new R-PHY session (1) or an existing R-PHY session (0). (See requirement 1 immediately following this table)
A	1 bit. Active bit. Indicates whether the R-PHY session is up (1) or down (0). (See requirement 2 below)

Requirements referenced in Table 11 above:

1. The Circuit Status AVP New bit SHOULD be set the first time the DEPI session is established after provisioning.
2. Regarding the Circuit Status AVP Active bit, once the CCAP Core is aware that the R-PHY session is down, the CCAP Core MUST NOT attempt to pass data traffic on the R-PHY session.

LCCEs utilize the Circuit Status AVP to indicate whether the R-PHY session is up and able to pass data traffic, or down and unable to pass data traffic. The Circuit Status ID does not control the RF output of the QAM Channel. Note that the Circuit Status AVP can be sent by both the CCAP Core and the RPD.

7.5.2 DEPI Control Specific AVP

This specification defines a single DEPI specific AVP for usage in control signaling. This is DEPI Pseudowire Subtype Capabilities List AVP.

7.5.2.1 DEPI Pseudowire Subtype Capabilities List AVP (SCCRQ, SCCRQ)

Figure 35 shows DEPI Pseudowire Subtype Capabilities List AVP. This AVP is used in the SCCRQ and SCCRQ messages to communicate which pseudowire subtypes are supported by the RPD and by the CCAP Core. There are two main PW Types defined for an RPD and a CCAP Core. These are DEPI MPT and DEPI PSP. This specification defines multiple variations or subtypes of the major pseudowire types. The DEPI Pseudowire Subtype Capabilities AVP provides the ability to communicate which pseudowire subtypes are supported by either device. The AVP is required to be present in SCCRQ and SCCRQ messages.




Figure 35 - DEPI Pseudowire Subtype Capabilities List AVP

Table 12 shows the set of defined values for pseudowire subtypes and their mapping to main PW types.

Table 12 - DEPI Pseudowire Subtypes

Pseudowire (PW) Subtype	Mnemonic	Value	PW Type
MPT DEPI Pseudowire Subtype	MPT-DEPI-PW	1	MPTPW
PSP Legacy-Pseudowire Subtype	PSP-LEGACY-PW	2	PSPPW
MCM Pseudowire Subtype	MCM-PW	3	MPTPW
PSP DEPI Multichannel Pseudowire Subtype	PSP-MULTICHAN-PW	4	PSPPW
reserved	reserved	5	N/A
PSP-UEPI-SCQAM Pseudowire Subtype	PSP-UEPI-SCQ	6	PSPPW
PSP-UEPI-OFDMA Pseudowire Subtype	PSP-UEPI-OFDMA	7	PSPPW
PSP-BW-REQ-SCQ Subtype	PSP-BW-REQ-SCQ	8	PSPPW
PSP-BW-REQ-OFDMA Subtype	PSP-BW-REQ-OFDMA	9	PSPPW
PSP-PROBE Subtype	PSP-BW-PROBE	10	PSPPW
PSP-RNG-REQ-SCQ Pseudowire Subtype	PSP-RNG-REQ-SCQ	11	PSPPW
PSP-RNG-REQ-OFDMA Pseudowire Subtype	PSP-RNG-REQ-OFDMA	12	PSPPW
PSP-MAP-SCQ Pseudowire Subtype	PSP-MAP-SCQ	13	PSPPW
PSP-MAP-OFDMA Pseudowire Subtype	PSP-MAP-OFDMA	14	PSPPW
PSP-SPECMAN Pseudowire Subtype	PSP-SPECMAN	15	PSPPW
PSP-PNM Pseudowire Subtype	PSP-PNM	16	PSPPW
Unused	reserved	17	N/A

Pseudowire (PW) Subtype	Mnemonic	Value	PW Type
MPT-55-1-RET Pseudowire Subtype	MPT-55-1-RET	18	MPTPW
MPT-55-2-FWD Pseudowire Subtype	MPT-55-2-FWD	19	MPTPW
MPT-55-2-RET Pseudowire Subtype	MPT-55-2-RET	20	MPTPW
PSP-NDF Pseudowire Subtype	PSP-NDF	21	PSPPW
PSP-NDR Pseudowire Subtype	PSP-NDR	22	PSPPW

The CCAP Core MUST indicate its support for pseudowire subtypes by including one or more of the PW Subtypes in the Pseudowire Subtype Capabilities List. The RPD MUST indicate its support for pseudowire subtypes by including one or more of the PW Subtypes in the Pseudowire Subtype Capabilities List.

7.5.3 DEPI Session Specific AVPs

The AVP types defined specifically for DEPI session signaling are shown in Table 13. The range of Attribute Types from 0-99 is reserved for DEPI Session Specific AVPs. These AVPs are only used in L2TP session messages. Note that several AVPs defined for DEPI have been deprecated from use in the R-PHY protocol. The CCAP Core MUST NOT transmit deprecated AVPs. The RPD MUST NOT transmit deprecated AVPs.

Table 13 - DEPI Defined Session Specific AVPs

Attribute Type	Description
0	Reserved
1	DEPI Result Code
2	DEPI Resource Allocation Request
3	DEPI Resource Allocation Reply
4	DEPI Local MTU
5	Deprecated (DOCSIS SYNC Control)
6	Deprecated (EQAM Capability Bits)
7	DEPI Remote MTU
8	DEPI Local UDP Port
9	Deprecated (DPR Session Type)
10	Deprecated (DPR Session Status)
11	DEPI Remote Multicast Join
12	DEPI Remote Multicast Leave
13	DEPI Multicast Capability
14	Upstream Flow
16	DEPI Pseudowire Subtype
17	DEPI L2-Specific Sublayer Subtype
18	DEPI DLM Response DSCP
19	DEPI S-BFD Return Path DSCP
20	DEPI Multicast UDP Port AVP

7.5.3.1 DEPI Result and Error Code (CDN, SLI)

Figure 36 shows the DEPI Result and Error Code AVP.




Figure 36 - DEPI Result and Error Code AVP

Table 14 - DEPI Result and Error Code AVP Details

Field	Description												
M	1 bit. Mandatory bit. (See Table Note).												
Length	10 bits. 8 bytes plus any additional bytes required for the Error Code and Error Message.												
Attribute Type	2 bytes. Set to 1.												
Result Code	2 bytes. Mandatory field. The values are as follows:												
	<table border="1"> <thead> <tr> <th>Result Code</th> <th>Result Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Session not established - Bad DEPI AVP Reference</td> </tr> <tr> <td>1</td> <td>Session not established - Bad PHY AVP Request</td> </tr> <tr> <td>2</td> <td>Session not established or disconnected for the reason indicated in Error Code.</td> </tr> </tbody> </table>	Result Code	Result Description	0	Session not established - Bad DEPI AVP Reference	1	Session not established - Bad PHY AVP Request	2	Session not established or disconnected for the reason indicated in Error Code.				
Result Code	Result Description												
0	Session not established - Bad DEPI AVP Reference												
1	Session not established - Bad PHY AVP Request												
2	Session not established or disconnected for the reason indicated in Error Code.												
Error Code	2 bytes. This field is optional. One cause of Error Code 0 is if the RPD does not have the values of M and N configured for the QAM Channel Symbol Rate. The values are as follows:												
	<table border="1"> <thead> <tr> <th>Error Code</th> <th>Error Description</th> </tr> </thead> <tbody> <tr> <td>0</td> <td>Device is not yet ready or configured properly.</td> </tr> <tr> <td>1</td> <td>Attempt to modify a locked PHY parameter.</td> </tr> <tr> <td>2</td> <td>Attempt to modify PHY Parameter failed - out of range value.</td> </tr> <tr> <td>3</td> <td>Requested PSP Flow PHB-IDs not supported.</td> </tr> <tr> <td>4</td> <td>Incorrect pseudowire type used in session.</td> </tr> </tbody> </table>	Error Code	Error Description	0	Device is not yet ready or configured properly.	1	Attempt to modify a locked PHY parameter.	2	Attempt to modify PHY Parameter failed - out of range value.	3	Requested PSP Flow PHB-IDs not supported.	4	Incorrect pseudowire type used in session.
Error Code	Error Description												
0	Device is not yet ready or configured properly.												
1	Attempt to modify a locked PHY parameter.												
2	Attempt to modify PHY Parameter failed - out of range value.												
3	Requested PSP Flow PHB-IDs not supported.												
4	Incorrect pseudowire type used in session.												
Error Message	Variable length. This field is optional.												

Table Note: The DEPI Result and Error Code AVP M field mandatory bit MUST be set to 0

The format of the fields of this AVP are identical to the corresponding fields in the standard L2TPv3 result and error code AVP except that the Vendor ID field has the value of 4491 rather than 0. The Result and Error Codes for this AVP are unique to [DEPI], and are in addition to the standard Result and Error codes listed in Section 7.5.1.2.

7.5.3.2 DEPI Resource Allocation Request (ICRQ)

Figure 37 shows the DEPI Resource Allocation Request AVP.




Figure 37 - DEPI Resource Allocation Request AVP

Table 15 - DEPI Resource Allocation Request AVP Details

Field	Description
M	1 bit. Mandatory bit. (See requirement 1 immediately following this table)
Length	10 bits. 6 bytes plus one additional byte for each flow that is requested.
Attribute Type	2 bytes. Set to 2.
PHB-ID	6 bits. Per Hop Behavior Identifier being requested by the CCAP Core for the RPD packet scheduler. Per hop behavior identifiers are defined in the "Per Hop Behavior Usage" section of [R-PHY]. When RPD supports direct mapping of downstream PSP flows to strict priority queues, PHB is set to 0.
Flow ID	3 bits. This is the flow ID requested by the CCAP Core. (See requirement 2 below)

Requirements referenced in table above:

1. The DEPI Resource Allocation Request AVP M field Mandatory bit MUST be set to a 1.
2. The CCAP Core MUST ensure that each flow ID is unique within a session.

A CCAP Core MUST include this AVP in an ICRQ for pseudowires with downstream sublayer data packets, as indicated in Table 10 and Table 24. This includes both PSP and non-PSP pseudowire types with downstream sublayer data. Each two bytes in the attribute payload represents a request for one unique downstream flow. Each request contains an assigned Flow ID and a corresponding PHB-ID that selects the per-hop behavior in the RPD, i.e., how it schedules packets from the flow relative to other flows in the same session (if any).

In a DEPI Resource Allocation Request AVP, the CCAP Core MAY request more than one downstream flow only for the pseudowire type PSP with subtypes PSP-MULTICHAN-PW and PSP-LEGACY-PW. For all other downstream pseudowire subtypes the CCAP Core MUST request only one flow with Flow ID equal to zero. In a DEPI Resource Allocation Request AVP with more than one downstream flow, the CCAP Core MAY request the same PHB-ID for more than one flow.

The RPD signals as GCP capabilities the maximum number of PSP flows it supports per PSP session. The CCAP Core MUST NOT request more PSP flows than the RPD advertises it is capable of supporting. The RPD MUST support Flow ID range from zero to (number of supported flows- 1). The CCAP Core MUST NOT assign Flow ID values that are outside of the range supported by the RPD. For example, if the RPD indicates support for a maximum of four PSP flows, the RPD needs to support and CCAP Core can only assign Flow ID values in the range 0-3.

The RPD signals as GCP capabilities the set of PHB-IDs it supports for its downstream packet scheduler and whether the RPD supports direct mapping of downstream PSP flows to strict priority queues. The CCAP Core MUST NOT request a PHB-ID for a downstream flow that is not advertised as supported by the RPD in a GCP capability. When the RPD supports direct mapping of downstream flows to strict priority queues, the PHB-ID signaling in DEPI Resource Allocation Request AVP is ignored.

If an RPD receives a DEPI Resource Allocation Request AVP with a PHB-ID that it has not advertised as supported, the RPD MUST tear down the session with CDN. The RPD MUST include in CDN message the Result Code 12 that indicates session DSCP mismatch as the reason for sending the CDN.

7.5.3.3 DEPI Resource Allocation Reply (ICRP)

This AVP is signaled by the RPD in an ICRP to inform the CCAP Core that it has accepted the flow IDs for downstream data packets of the session, as requested by the CCAP Core in a DEPI Resource Allocation Request.

Figure 38 shows the DEPI Resource Allocation Reply AVP.




Figure 38 - DEPI Resource Allocation Reply AVP

Table 16 - DEPI Resource Allocation Reply AVP Details

Field	Description
M	1 bit. Mandatory bit. (See requirement immediately following this table.)
Length	10 bits. 6 bytes plus two additional bytes for each flow that is acknowledged.
Attribute Type	2 bytes. Set to 3.
PHB-ID	6 bits. Per Hop Behavior Identifier being accepted by the RPD. Per hop behavior identifiers are defined in the "Per Hop Behavior Usage" section of [R-PHY].
Flow ID	3 bits. This is the flow ID accepted by the RPD. The flow ID is unique within a session.

The DEPI Resource Allocation Reply AVP M field Mandatory bit MUST be set to a 1.

A CCAP Core MUST tear down a session for which the RPD does not reply with all requested flows in its DEPI Resource Allocation Reply AVP. Note that this operation differs from that described for M-CMTS in [DEPI], which permitted partial allocation of requested downstream flows.

7.5.3.4 DEPI Local MTU (ICRQ)

Figure 39 shows the DEPI Local MTU AVP.




Figure 39 - DEPI Local MTU AVP

Table 17 - DEPI Local MTU AVP Details

Field	Description
M	1 bit. Mandatory bit. (See requirement immediately following this table.)
Length	10 bits. Set to 8.
Attribute Type	2 bytes. Set to 4.
DEPI Local MTU	2 bytes. In the ICRQ message, this is the MTU (Maximum Transmission Unit) that the CCAP Core can receive from the RPD on the CIN interface.

The DEPI Local MTU AVP M field Mandatory bit MUST be set to a 1.

The MTU is the Layer 3 payload of a Layer 2 frame. For R-DEPI, the MTU includes the L2TPv3 header and payload and the IP header, but does not include the Ethernet Header or the CRC. For example, a 1518-byte Ethernet frame (1522 bytes if VLAN tags are present) would support an MTU of 1500 bytes.

7.5.3.5 DEPI Remote MTU Max Payload (ICRP)

Figure 40 shows the DEPI Remote MTU Max Payload AVP.




Figure 40 - DEPI Remote MTU Max Payload AVP

Table 18 - DEPI Remote MTU Max Payload AVP Details

Field	Description
M	1 bit Mandatory bit. (See requirement immediately following this table.)
Length	10 bits. Set to 8.
Attribute Type	2 bytes. Set to 7.
DEPI Remote MTU	2 bytes. In the ICRP message, this is the MTU that the RPD can receive from the CCAP Core on the CIN interface. The MTU is the Layer 3 payload of a Layer 2 frame.

The sender (RPD) MUST set the M field Mandatory bit to 1 in DEPI Remote MTU Max Payload AVP.

7.5.3.6 DEPI Local UDP Port (ICRQ, ICRP)

Figure 41 shows the DEPI Local UDP Port AVP.




Figure 41 - DEPI Local UDP Port AVP

Table 19 - DEPI Local UDP Port AVP Details

M	1 bit. The Mandatory bit MUST be set to a 1.
Length	10 bits. Set to 8.
Attribute Type	2 bytes. Set to 8.
Local UDP Port	16 bits. UDP port to be used for session packets that are being sent to the local LCP.

The CCAP Core would issue this AVP during the session setup (ICRQ) if UDP is enabled and if the CCAP Core wants a data session to use a different UDP port for sending session packets from the RPD to the CCAP Core other than the UDP port that was negotiated during the Control Connection setup. Similarly, the RPD would issue this AVP in ICRP message if UDP is enabled and if the RPD wants a data session to use a different UDP port for sending session packets from the CCAP Core to the RPD other than the UDP port that was negotiated during the Control Connection setup.

The CCAP Core MUST support the Local UDP Port AVP. The RPD MAY support the Local UDP Port AVP. The RPD MUST NOT use this AVP in setup of multicast sessions.

7.5.3.7 DEPI Remote Multicast Join

Figure 42 shows the DEPI Remote Multicast Join AVP.




Figure 42 - DEPI Remote Multicast Join AVP

Table 20 - DEPI Remote Multicast Join AVP Details

Field	Description
M	1 bit Mandatory bit. (See requirement immediately following this table.)
Length	10 bits. Set to 40.
Attribute Type	2 bytes. Set to 11.
Source IP Address	16 bytes. This field signals the Source IP address of the multicast stream the RPD should join. IPv4 address is encoded on the first four bytes of this field.
Group IP Address	16 bytes. This field signals the Group (Destination) IP address of the multicast stream the RPD should join. IPv4 address is encoded on the first four bytes of this field.

The DEPI Remote Multicast Join AVP M field Mandatory bit MUST be set to a 1.

7.5.3.8 DEPI Remote Multicast Leave

Figure 43 shows the DEPI Remote Multicast Leave AVP.




Figure 43 - DEPI Remote Multicast Leave AVP

Table 21 - DEPI Remote Multicast Leave AVP Details

Field	Description
M	1 bit Mandatory bit. (See requirement immediately following this table.)
Length	10 bits. Set to 40.
Attribute Type	2 bytes. Set to 12.
Source IP Address	16 bytes. This field signals the Source IP address of the multicast stream the RPD should leave. IPv4 address is encoded on the first four bytes of this field
Group IP Address	16 bytes. This field signals the Group (Destination) IP address of the multicast stream the RPD should leave. IPv4 address is encoded on the first four bytes of this field.

The DEPI Remote Multicast Leave AVP M field Mandatory bit MUST be set to a 1.

The CCAP Core SHOULD signal the same Source and Group IP address in the DEPI Remote Multicast Leave AVP as was signaled in the DEPI Remote Multicast Join AVP for this session. In case the RPD receives DEPI Remote Multicast Leave AVP with Source and/or Group IP address that do not match the address(es) sent via the DEPI Remote Multicast Join AVP, the RPD SHOULD handle it as though a Leave was received for the multicast stream associated with that session.

7.5.3.9 DEPI Multicast Capability (SCCRQ, SCCRP)

Figure 44 shows the DEPI Multicast Capability AVP.




Figure 44 - DEPI Multicast Capability AVP

Table 22 - DEPI Multicast Capability AVP Details

Field	Description
M	1 bit. Mandatory bit. (See requirement immediately following this table.)
Length	10 bits. Set to 8.
Attribute Type	2 bytes. Set to 13.
C	DEPI Multicast Capability. When set to '1' the sender of the AVP supports DEPI Multicast signaling on downstream sessions.

The DEPI Remote Multicast Leave AVP M field Mandatory bit MUST be set to a 1.

DEPI Multicast Capability AVP indicates the ability of the CCAP Core or of the RPD to support the multicast operation, including the signaling of multicast session Ids assigned by the CCAP Core.

The CCAP Core MUST support the DEPI Multicast Capability AVP. The RPD MUST support the DEPI Multicast Capability AVP.

The CCAP Core SHOULD set the value of the C bit to "1" in the DEPI Multicast Capability AVP. The RPD MUST set the value of the C bit to "1" in the DEPI Multicast Capability AVP.

7.5.3.10 Upstream Flow (ICRQ)

The Upstream Flow AVP is signaled by the CCAP Core in an ICRQ to inform the RPD of Flow Ids and DSCPs of data packets for each upstream flow in the session.

Figure 45 shows the Upstream Flow AVP.




Figure 45 - Upstream Flow AVP

Table 23 - Upstream Flow AVP Details

Field	Description
M	1 bit. Mandatory bit. (See requirement immediately following this table.)
Length	10 bits. 6 bytes plus 2 additional bytes for each flow that is described.
Attribute Type	2 bytes. Set to 14.
PHBID	6 bits. Per Hop Behavior Identifier that equals the 6-bit DSCP with which the RPD transmits L2TPv3 data packets for the flow.
Flow ID	3 bits. This is an upstream flow ID assigned by the CCAP Core. The flow ID is unique within a session. For pseudowires operating with a single flow, such as pseudowires which do not use PSP Sublayer or PSP pseudowires for which the configuration of multiple flows is prohibited, the CCAP Core uses Flow ID value of 0.

The Upstream Flow AVP M field Mandatory bit MUST be set to a 1.

The CCAP Core MAY include more than one flow id for pseudowires with PSP-UEPI-SCQAM and PSP-UEPI-OFDMA L2-Specific Sublayer Subtypes in ICRQ message. For pseudowires with L2-Specific Sublayer Subtypes other than PSP-UEPI-SCQAM and PSP-UEPI-OFDM the CCAP CORE MUST include exactly one Upstream Flow AVP with Flow ID equal to zero in ICRQ message.

7.5.3.11 DEPI Pseudowire Subtype AVP (ICRQ)

Figure 46 shows the DEPI Pseudowire Subtype AVP.




Figure 46 - DEPI Pseudowire Subtype AVP

R-DEPI uses the Pseudowire Subtype values defined in Table 12 to indicate the PW subtype of DEPI session requested.

7.5.3.12 DEPI L2-Specific Sublayer Subtype AVP (ICRQ, ICRP, ICCN)

Figure 47 shows the DEPI L2-Specific Sublayer Subtype AVP. Table 24 defines the set of valid values for the DEPI L2-Specific Sublayer Subtype.




Figure 47 - DEPI L2-Specific Sublayer Subtype AVP

Table 24 - DEPI L2-Specific Sublayer Subtype Values

L2-Specific Sublayer Subtype	Value	Direction	L2-Specific PW Sublayer Type
MPT DEPI L2-Specific Sublayer Subtype	1	Downstream	MPT
PSP Legacy L2-Specific Sublayer Subtype	2	Downstream	PSP
MCM Specific L2-Sublayer Subtype	3	Downstream	MPT
PSP DEPI Multichannel L2-Specific Sublayer Subtype	4	Downstream	PSP
reserved	5	n/a	n/a
PSP-UEPI-SCQAM L2-Specific Sublayer Subtype	6	Upstream	PSP
PSP-UEPI-OFDMA L2-Specific Sublayer Subtype	7	Upstream	PSP
PSP-BW-REQ-SCQ L2-Specific Sublayer Subtype	8	Upstream	PSP
PSP-BW-REQ-OFDMA L2-Specific Sublayer Subtype	9	Upstream	PSP
PSP-PROBE L2-Specific Sublayer Subtype	10	Upstream	PSP
PSP-RNG-REQ-SCQ L2-Specific Sublayer Subtype	11	Upstream	PSP
PSP-RNG-REQ-OFDMA L2-Specific Sublayer Subtype	12	Upstream	PSP
PSP-MAP-SCQ L2-Specific Sublayer Subtype	13	Downstream	PSP
PSP-MAP-OFDMA L2-Specific Sublayer Subtype	14	Downstream	PSP
PSP-SPECMAN L2-Specific Sublayer Subtype	15	Upstream	PSP
PSP-PNM L2-Specific Sublayer Subtype	16	Upstream	PSP
reserved	17	n/a	n/a
MPT-55-1-RET L2-Specific Sublayer Subtype	18	Upstream	MPT
MPT-55-2-FWD L2-Specific Sublayer Subtype	19	Downstream	MPT
MPT-55-2-RET L2-Specific Sublayer Subtype	20	Upstream	MPT
PSP-NDF L2-Specific Sublayer Subtype	21	Downstream	PSP
PSP-NDR L2-Specific Sublayer Subtype	22	Upstream	PSP

The CCAP Core MUST include the L2-Specific Sublayer AVP in the ICRQ, ICRP and ICCN messages, indicating the L2-Specific Sublayer Header Subtype consistent with the pseudowire type for the DEPI flow. The RPD MUST include the L2-Specific Sublayer AVP in the ICRP message, indicating the L2-Specific Sublayer Header Type consistent with the Pseudowire type for the R-PHY protocol flow.

7.5.3.13 DEPI DLM Response DSCP AVP (ICRQ)

Figure 48 shows the DEPI DLM Response DSCP AVP.




Figure 48 - DEPI DLM Response DSCP AVP

Table 25 - DEPI DLM Response DSCP AVP Details

M	1 bit. The Mandatory bit MUST be set to a 1.
Length	10 bits. Set to 8.
Attribute Type	2 bytes. Set to 18.
PHB Code	The PHB Code encoding is defined by [RFC 3140]

The CCAP Core MAY include DEPI DLM Response DSCP AVP in ICRQ message. Additional requirements for the use of this AVP are described in Section 7.3.3.3.2.

7.5.3.14 DEPI S-BFD Return Path DSCP AVP (ICRQ)

Figure 49 shows the DEPI S-BFD Return Path DSCP AVP.

**Figure 49 - DEPI S-BFD Return Path DSCP AVP****Table 26 - DEPI S-BFD Return Path DSCP AVP Details**

M	1 bit. The Mandatory bit MUST be set to a 1.
Length	10 bits. Set to 8.
Attribute Type	2 bytes. Set to 19.
PHB Code	The PHB Code encoding is defined by [RFC 3140]

The CCAP Core MAY include DEPI S-BFD Return Path DSCP AVP in ICRQ message. Additional requirements for the use of this AVP are described in section 7.3.3.3.2.

7.5.3.15 DEPI Multicast UDP Port (ICRQ, ICRP)

Figure 50 shows the DEPI Multicast UDP Port AVP.

**Figure 50 - DEPI Multicast UDP Port AVP****Table 27 - DEPI Multicast UDP Port AVP Details**

M	1 bit. The Mandatory bit MUST be set to a 1.
Length	10 bits. Set to 8.
Attribute Type	2 bytes. Set to 20.
Multicast UDP Port	16 bits. UDP port to be used for multicast session packets that are being sent by the CCAP Core.

The CCAP Core would issue this AVP during the session setup (ICRQ), if UDP is enabled and if the CCAP Core wants a data session to use a different UDP port for sending multicast session packets from the CCAP Core to the RPD, other than the UDP port that was negotiated during the Control Connection setup.

The CCAP Core MUST support the DEPI Multicast Port AVP. The RPD MAY support the DEPI Multicast UDP Port AVP. The QAM Channel Physical Layer specific AVP types initially defined in [DEPI] have been deprecated from use in the R-PHY protocol.

8 DEPI FORWARDING PLANE

8.1 L2TPv3 Transport Packet Format

8.1.1 Data Message Format without UDP

Figure 51 shows the format of the L2TPv3 data packet outer encapsulation (without UDP headers).




Figure 51 - L2TPv3 Data Packet Outer Encapsulation (without UDP)

8.1.2 Data Message with a UDP Header

Figure 52 shows the format of the L2TPv3 data packet with a UDP Header.




Figure 52 - L2TPv3 Data Packet Outer Encapsulation with UDP Header

8.1.3 Specific Headers for Data Messages

8.1.3.1 L2TPv3 Data Header

The L2TPv3 Data Header is not used in R-DEPI because data packets with L2TPv3 over IP do not require an L2TPv3 Data Header.

8.1.3.2 L2TPv3 DEPI Sublayer Header

DEPI supports three main pseudowire types. The first type, known as D-MPT, is used for transporting MPEG packets. The second type, known as MCM, is a variation of the D-MPT type. MCM allows transport of MPEG packet belonging to multiple QAM channels in a single DEPI packet. The third type, known as PSP, is used for transporting DOCSIS frames. Each pseudowire type has a unique L2TPv3 DEPI sublayer header format. The fields of these sublayer headers are defined in Sections 8.2 and 8.3. Additionally, both pseudowire types support a latency measurement sublayer header, the fields of which are defined in Section 8.4.1.

The CCAP Core MUST use the appropriate DEPI sublayer header for the pseudowire type of the DEPI session.

The RPD MUST accept packets from the CCAP Core that contain the appropriate DEPI sublayer header for the negotiated pseudowire type. The RPD MUST send a CDN message to tear down a session in which packets are received with the wrong pseudowire type. The RPD MUST ignore packets received that do not comply with these L2TPv3 DEPI sublayer header definitions.

8.1.3.3 DEPI Payload

The payload contains one or more segments. In D-MPT mode and in MCM mode, each segment is a 188 byte MPEG packet. In PSP mode, a segment contains either a full DOCSIS frame or a partial DOCSIS frame or data intended for the PLC channel.

8.2 DEPI MPT Sublayer Header

Figure 53 shows the format of the DEPI MPT Sublayer header and payload.




Figure 53 - DEPI MPT Sublayer Header and Payload

Table 28 - Header Format for MPT Sublayer Header and Payload Details

Field	Description
V	1 bit. VCCV bit. Set to 0. Reserved for compatibility with [RFC 5085].
S	1 bit. Sequence bit. Set to 1 to indicate that the sequence number field is valid. Set to 0 to indicate that the sequence field is not valid.
H	2 bits. Extended Header bits. Set to '00' to indicate a DEPI sublayer header that matches the current active pseudowire type (either D-MPT or PSP).
X	1 bit. Reserved bit. Set to all zeroes at the transmitter, ignored at the receiver.
Flow ID	3 bits. Flow Identifier.
Reserved	1 byte. Reserved field. Set to all zeroes at the transmitter, ignored at the receiver.

Field	Description
Seq Num	2 bytes. Sequence Number. The sequence number increments by one for each data packet sent, and may be used by the receiver to detect packet loss. (See requirement following this table)

The initial value of the Sequence Number SHOULD be random (unpredictable).

The CCAP Core MUST support the creation of a single DEPI MPT session per SC-QAM channel. The RPD MUST support a single DEPI MPT session per SC-QAM channel. The CCAP Core MUST NOT put null bytes between the UDP header and the first MPEG-TS header or between consecutive MPEG-TS packets. The CCAP Core MUST support all bits in the MPT sublayer header.

The RPD MUST accept 1-7 MPEG-TS packets in an L2TPv3 payload when the path MTU is 1500 bytes in length. The length of an Ethernet frame containing seven MPEG-TS packets with L2TPv3 with a D-MPT L2TPv3 sublayer, and an IPv4, [IEEE 802.1q] header is 1378 bytes. If the RPD, the CCAP Core, and the network between them all allow larger MTU sizes, the CCAP Core MAY increase the total number of MPEG-TS packets transmitted per L2TPv3 packet.

The CCAP Core MAY insert null MPEG packets into the D-MPT or MCM stream for a DOCSIS SC-QAM channel. A null MPEG packet is 188 bytes in length with a reserved PID value of 0x1FFF as defined in [ISO 13818-1]. The RPD MAY discard these null MPEG packets. The RPD is only required to support one flow for D-MPT or MCM mode.

8.3 DEPI Multi-channel MPEG (MCM) L2TPv3 Sublayer Header Format

Table 29 - Sublayer "H" Bits Decode Table

Original DEPI vs. DEPI-MCM	Sublayer H Bits	Direct Decode of "H" Bits	L2TPv3 Control Plane Control
Orig DEPI	2'B00	DEPI sublayer header that matches the current active pseudowire type (either D-MPT or PSP).	L2TPv3 Control Plane sets up pseudowire choice of D-MPT or PSP.
Add-on to Orig DEPI	same = 2'B00	Reuses existing DEPI H bits of 2'B00.	Downstream PHY Module In-Band Control Plane selects new MCM Pseudowire type.
Orig DEPI	2'B01	DEPI Latency Measurement (DLM) Sublayer Header	No control over DLM packets.
Orig DEPI	2'B10 & 2'B11	Not defined	NA

The MCM L2TPv3 Sublayer Header is shown in Figure 54. The MCM Sublayer Header carries multiple channels inside an IP frame (unlike the DEPI D-MPT mode where each IP frame carried the same channel). The MCM Sublayer packet takes advantage of recent industry changes to carry 2000 byte Ethernet frames. This allows the MCM DEPI packet to carry up to 10 MPEG-TS packets.

The "H" bits in the Sublayer first double word partially define the type of sublayer header. In D-MPT, PSP, and MCM mode, the Control Plane sets bits to identify the packet content.




Figure 54 - MCM Sublayer Header and Payload

Table 30 - L2TPv3 MCM Sublayer Header and Payload Format

Field	Description
V	1 bit. VCCV bit. Set to 0. Reserved for compatibility with [RFC 5085].
S	1 bit. Sequence bit. Set to 1 to indicate that the sequence number field is valid. Set to 0 to indicate that the sequence field is not valid.
H	2 bits. Extended Header bits. Set to '00' exactly the same as existing DEPI; the In-Band Ethernet Control Plane chooses MCM Sublayer to indicate an MCM sublayer header is being used to move multiple channels of MPEG in the same IP Frame.
X	1 bit. Reserved bit. Set to all zeroes at the transmitter, ignored at the receiver.
Flow ID	3 bits. Flow Identifier.
Reserved	4 bits. Reserved field. Set to all zeroes at the transmitter, ignored at the receiver.
Packet Count	4 bits. Packet Count is the number of MPEG-TS packets in the L2TPv3 payload. Set to 4'h0=reserved, 4'h1=1pkt, 4'h2=2pkts...4'hA=10pkts, 4'hB>4'hF=reserved. Packet Count also indicates which MPTS TAGs and MPTS Sequence Number fields are valid.

Field	Description
Seq Num	2 bytes. Sequence Number. The sequence number increments by one for each data packet sent, and may be used by the receiver to detect packet loss. The sequence number can be used to track lost IP frames. (See requirement immediately following this table)
MPTS TAG	1 byte. MPTS Tag. The MPTS Tag provides a pointer to the RPD's MPEG-TS Stream Processing Resource (SPR). The MCM IP frame carries multiple MPEG-TS packets all directed at different MPEG-TS processing engines and QAM channels. The value of the Packet Count field defines how many MPTS tags are valid.
MPTS Sequence Number	4 bits. MPTS Sequence Number. The MPTS Sequence Number provides per-channel tracking information for the RPD. The RPD receives MPTS TAGs and MPTS Sequence Numbers, these two entities allow the RPD to count per-channel the MPEG-TS packets as they enter the RPD. The RPD will have the capability to replace per-channel MPEG-TS packets with MPEG NULL packets.

In the L2TPv3 MCM Sublayer Header, the initial value of the Sequence Number SHOULD be random (unpredictable).

The new L2TPv3 MCM sublayer header has been defined so the R-PHY protocol can carry multiple channels of data in the same IP frame. It consists of a 4-byte-long subheader and 20-byte-long MPTS TAG and Sequence Number array. This array contains attributes associated with MPEG frames carried in the L2TPv3 payload. The array is fixed in size and supports up to 10 MPEG-TS packets. The value of the Packet Count field indicates which entries in the array are valid. For example, if the Packet Count field contains a value of 8, the first eight entries in the array are valid. The remaining, invalid entries in the array, are transmitted as zeroes and ignored upon reception.

The L2TPv3 Sequence Number is helpful for tracking possible lost and/or out-of-order IP frames. The MPTS Sequence Number helps in dynamically correcting channel-specific MPEG packets. The combination of the L2TPv3 Sequence Number and the MPTS Sequence Number allow for error-recovery accounting in the event an IP frame with varied channel-specific data is lost.

8.4 PSP Sublayer Header

8.4.1 PSP Legacy Sublayer Header

Figure 55 shows the format of the DEPI Legacy PSP Sublayer header and payload which was defined in the [DEPI] specification.




Figure 55 - DEPI Legacy PSP Sublayer Header and Payload

Table 31 - DEPI Legacy PSP Sublayer Header and Payload Format Details

Field	Description
V	1 bit. VCCV bit. Set to 0. Reserved for compatibility with [RFC 5085].
S	1 bit. Sequence bit. Set to 1 to indicate that the sequence number field is valid. Set to 0 to indicate that the sequence field is not valid.
H	2 bits. Extended Header bits. Set to '00' to indicate a DEPI sublayer header which matches the current active pseudowire type (either D-MPT or PSP).
X	1 bit. Reserved bit. Set to all zeros at the transmitter, ignored at the receiver.
Flow ID	3 bits. Flow Identifier.
Segment Count	7 bits. This is the number of segments in the DEPI PSP Payload, and this is also the number of 2 byte entries in the PSP Segment Table.
Seq Num	2 bytes. Sequence Number. The sequence number increments by one for each data packet sent, and may be used by the receiver to detect packet loss. In the first packet, the sender inserts random (unpredictable) initial value of Sequence Number.
B	Begin bit. 1 bit. Set to a 1 to indicate that the PSP Frame contains the beginning of a DOCSIS frame. Otherwise, set to 0.
E	End bit. 1 bit. Set to a 1 to indicate that the PSP Frame contains the end of a DOCSIS frame. Otherwise, set to 0.
Segment Length	14 bits. Length of DEPI segment in bytes.

Neither RPD nor the CCAP Core uses the legacy DEPI PSP Sublayer header format. The legacy DEPI PSP Sublayer header format is provided in this document for historical reference purposes only. Additional details can be found in [DEPI].

8.4.2 PSP Sublayer Header

The PSP can take a series of DOCSIS frames, assemble them into a stream of back-to-back DOCSIS frames, and then split that stream up into PSP PDUs. In doing so, the first and last DOCSIS frames of a PSP PDU may be split into segments across PSP PDUs. DOCSIS frames that are not the first or last frame in a PSP PDU will not be split.

A DOCSIS frame may be split into more than two segments and therefore may be spread across more than two PSP PDUs. PSP segments can also carry Message Blocks for the PLC sub-channel of the OFDM channel. This specification defines two types of segments DOCSIS PDU Segments and PLC Message Block Segments. While PSP supports segmentation of DOCSIS PDUs, PLC Message Blocks cannot be segmented.

The Segment Table provides information on the contents of each of the subsequent PSP frames. This includes signifying if the frame is the beginning, middle, end, or comprises an entire DOCSIS frame.

The main format difference between the DEPI Legacy PSP Sublayer header and the PSP Sublayer header defined in this specification is the expansion of the segment table from 16-bits to 32-bits. 32-bit long entries of the segment table carry additional information such as channel identifier, per channel sequence number, and OFDM profile information.

Figure 56 shows the unified format of the DEPI PSP Sublayer header that is utilized for carrying data intended for downstream SC-QAM channels, OFDM channels, and the PLC sub-channels of the OFDM channels.




Figure 56 - DEPI PSP Sublayer Header and Payload

Table 32 - DEPI PSP Sublayer Header and Payload Format Details

Field	Description
V	1 bit. VCCV bit. Set to 0. Reserved for compatibility with [RFC 5085].
S	1 bit. Sequence bit. Set to 1 to indicate that the sequence number field is valid. Set to 0 to indicate that the sequence field is not valid.
H	2 bits. Extended Header bits. Set to '00' to indicate a DEPI sublayer header that matches the current active pseudowire type.
Flow ID	3 bits. Flow ID. Allows for multiple flows per session.
X	1 bit. Reserved bit. Set to all zeroes at the transmitter, ignored at the receiver.
Segment Count	7 bits. This is the number of segments in the DEPI PSP payload, and this is also the number of 4 byte entries in the PSP segment table.
Seq Num	2 bytes. Sequence Number. (See requirements 1a and 1b immediately following this table)
B	Begin bit. 1 bit. The interpretation of the "B" bit depends on the type of the segment. For DOCSIS PDU Segments, the "B" bit is set to 1 to indicate that the PSP segment contains the beginning of a DOCSIS frame. Otherwise, set to 0. For PLC Message Block Segments, the value of the "B" bit indicates the presence of the prepended 32-bit timestamp to the message block located within the segment payload. In such case (if the "B" bit is set to 1), then the PLC Message Block in the segment is prepended with a 32-bit timestamp. If the "B" bit is set to 0, the segment contains a PLC Message Block and no timestamp.
E	End bit. 1 bit. For DOCSIS PDU Segments the "E" bit is set to 1 to indicate that the PSP frame contains the end of a DOCSIS frame. Otherwise, set to 0. For PLC Message Block Segments the usage of the "E" bit is undefined. The sender transmits the "E" bit as 0, and the receiver ignores it.
Segment Length	14 bits. Length of DEPI segment in bytes.

Field	Description
Channel ID	8 bits. An identifier of the downstream channel or the PLC sub-channel to which a particular segment belongs. The CCAP Core provides the mapping between downstream channels in the RPD and Channel IDs used in the segment table via the L2TPv3 control plane protocol. See requirement 2 immediately following this table. Note that while the size of the Channel ID field permits effective identification and multiplexing of data from up to 256 channels, this specification imposes additional restrictions on how many and which channels can be grouped in a single session. Additional detail can be found in Section 8.4.3.2. The CCAP Core communicates the mapping between Channel ID field and the Remote End Id AVP.
Channel Seq. Num	4 bits. Per channel sequence number. The sequence number is incremented by the sender (CCAP Core) for each segment of a particular channel on a particular PSP flow. This field can be used by the receiver (the RPD) in the reassembly process to detect lost segments.
Profile ID	4 bits. The interpretation of the Profile ID field depends on the type of the segment and the channel type it is associated with. For DOCSIS PDU Segments carrying data intended for SC-QAM channels, the sender transmits this field as "0000". This field is ignored by the receiver. For DOCSIS PDU Segments carrying data for OFDM channels, this field indicates the OFDM Profile ID of a downstream packet. The Profile ID field is valid in the first segment of a DOCSIS PDU (if the "B" bit is set to 1). If the "B" bit is set to 0, the sender transmits the profile field as '0000' and the receiver ignores it. For segments associated with an OFDM PLC sub-channel, the lowest bit of this field is referred to as the Segment Type or "T" bit. The three most significant bits of this field are transmitted as '000' by the CCAP Core and ignored by the RPD. If the "T" bit is set to 0 this indicates a DOCSIS PDU Segment type. If the "T" bit is set to 1 this indicates a PLC Message Block Segment type.

Requirement referenced in Table above:

- 1a. The Sequence Number in the DEPI PSP Sublayer Header MAY be used by the receiver to detect packet loss.
- 1b. The initial value of the Sequence Number in the DEPI PSP Sublayer Header SHOULD be random (unpredictable).
2. When sending PSP messages, the sender MUST set the Channel ID in every entry of the segment table in the DEPI PSP Sublayer Header.

The CCAP Core MUST increment the sequence number by one for each data packet sent on a particular PSP flow.

The CCAP Core (the transmitter) SHOULD insert a random (unpredictable) initial value of the sequence number.

The RPD (the receiver) MAY use the sequence number to detect packet loss.

The CCAP Core MUST increment the channel sequence number by one for each sent segment for a particular channel on a particular PSP flow.

The RPD MAY use the channel sequence number to detect packet loss on a particular channel and a particular PSP flow.

When transmitting a DOCSIS PDU for OFDM data channels, the CCAP Core MUST specify the OFDM profile ID in the first segment of the PDU.

When transmitting a DOCSIS PDU for OFDM data channels, the CCAP Core MUST set the Profile ID to '000' in all but the first segment of the PDU.

When transmitting a DOCSIS PDU for SC-QAM data channels, the CCAP Core MUST set the Profile ID to '0000' in all segments.

8.4.3 Interpretation of PSP Sublayer Header for PLC Sub-channel

Figure 57 shows an example of the DEPI PSP Sublayer header and payload with segments intended for a PLC sub-channel. The PSP header format allows the transport of DOCSIS PDUs for transmission on the message channel of the PLC as well as for the transport of PLC Message Blocks.




Figure 57 - Use of DEPI PSP Sublayer Header for PLC Sub-channel

The Profile ID field in the segment header includes a Segment Type "T" bit that allows the receiver to differentiate segments carrying DOCSIS PDU segments from segments carrying PLC Message Blocks. When the "T" bit is set to 1, the segment carries a complete PLC Message Block including a space for a 24-bit CRC message block. The CCAP Core does not initialize the CRC field in the message block and typically fills it with null bytes. The RPD is responsible for computing and insertion of the 24-bit CRC. Any of the transported message blocks can be preceded by a 32-bit DOCSIS timestamp. The presence of the timestamp is indicated when both the "T" bit is set to 1 and the "B" bit is set to 1 in the segment table.

8.4.3.1 Rules for Transport of PLC Data

A single PSP session and even a single PSP flow can be used to transport a mix of DOCSIS PDUs for the OFDM data channel and the PLC sub-channel as well as PLC message blocks for a single OFDM channel. The CCAP Core is bound by the requirements of [MULPIv3.1] in the selection of DOCSIS message types that can be sent on the PLC sub-channel.

The CCAP Core can use PSP sessions to transport Trigger and EM Message Blocks as well as any other message blocks that will be defined in the future. Timestamp message blocks are generated locally by the RPD.

When the 32-bit Timestamp field is prepended to a PLC Message Block, its value communicates to the downstream RPD the time when the message block is to be transmitted. The CCAP Core MUST insert the timestamp value with enough accuracy to uniquely identify the PLC frame in which the message block will be sent. The CCAP Core MUST send any message blocks that are preceded by the timestamp with time advance equivalent to the duration of 1-20 PLC frames to account for DEPI transmission path delays and necessary processing time in the downstream RPD. An RPD MAY discard message blocks with a timestamp outside of the advance time window equivalent to 1-20 PLC frame durations.

The CCAP Core MUST send Trigger Message Blocks in the order in which they are to be transmitted. Trigger Message Blocks MUST be sent with a prepended timestamp. When the timestamp is not prepended, the downstream RPD autonomously decides into which PLC frame to insert the Trigger message block.

The CCAP Core MUST transmit EM Message Blocks in the order in which they are to be transmitted on the PLC. The CCAP Core MUST send all EM Message Blocks with a prepended timestamp. If the CCAP Core sends more than one EM Message Block scheduled for transmission in the same PLC frame, the downstream RPD MUST transmit EM Message Blocks in the same order in which they have been received.

[MULPIv3.1] defines a generic format to be used for Message Block types 5-15, with the details of these message blocks to be defined in the future. The CCAP Core MUST NOT send a Message Block type 5-15 with a Message Body Size larger than 343 (total Message Block size of 348 bytes) if the Message Block is to be transmitted on the PLC of an OFDM channel with a 4K FFT size. Such a Message Block is too large to fit in a single PLC frame and thus can never be transmitted. Message Block types 5-15 can be sent either with or without a prepended timestamp.

The CCAP Core MUST limit the volume of data sent on the PLC pseudowire to match the traffic rate supported by the PLC. This requirement is defined to prevent excessive buffering on the downstream RPD. Additionally, when sending Message Blocks with prepended timestamps, the CCAP Core MUST ensure that the total size of Message Blocks assigned to a given PLC frame does not exceed the capacity of that frame after accounting for the Timestamp Message Block. [MULPIv3.1] gives PLC frame capacities prior to accounting for the Timestamp Message Block.

8.4.3.2 PSP Requirements and Use Rules

The RPD MUST support signaling and operation with at minimum four PSP flows per session. The RPD MAY support up to eight PSP flows per session.

The PSP pseudowire header format includes an 8 bit Channel ID field, which allows effective identification and multiplexing of up to 256 channels or PLC sub-channels in a single session. Requiring support for arbitrary combinations of channels in a single session would complicate the implementation of the RPD. For this reason this specification only mandates support for two combinations of channel types in a single L2TPv3 session. The support for other channel type combinations is not expressly prohibited. Their definition and methods of signaling are outside of the scope of this specification.

The RPD MUST support a single PSP session with 1-32 SC-QAM channels mapped to the same downstream RF port of the RPD.

The CCAP Core MUST support a single PSP session with one SC-QAM channel.

The CCAP Core MAY support a single PSP session with more than one SC-QAM channel mapped to the same downstream RF port of the RPD.

The RPD MUST support a single PSP session with exactly one OFDM data channel and one PLC sub-channel of the same OFDM channel.

The CCAP Core MUST support a single PSP session with exactly one OFDM data channel and one PLC sub-channel of the same OFDM channel.

To minimize DEPI overhead; a single PSP frame can transport multiple DOCSIS frames with different OFDM profiles. For example, a single DEPI packet can include a segment with a DOCSIS MDD message intended for transmission on profile "A" and a segment with a DOCSIS data PDU intended for transmission on profile "C". Similarly, a single PSP frame can transport multiple DOCSIS frames with different Channel IDs.

PSP pseudowires are intended to be implemented with a single segmentation and reassembly (SAR) context for each PSP flow. Within a single PSP flow, the CCAP Core MUST complete the transmission of all PSP segments of a DOCSIS frame before transmitting the first segment of another DOCSIS PDU even if it belongs to a different OFDM profile or a different channel.

8.5 Optional Headers

The DEPI protocol supports optional VLAN [IEEE 802.1q] and MPLS headers. An example of a complete DEPI packet, including VLAN and MPLS headers, is shown in Figure 58.




Figure 58 - Example of a Complete DEPI Packet Format (IPv4) with Extras

The specific requirements for support of VLAN and MPLS headers by the CCAP Core and the RPD are to be determined.

8.6 DEPI Latency Measurement (DLM) Sublayer Header

The DEPI Latency Measurement (DLM) Sublayer Header is used by the CCAP Core and the RPD to measure the delay and latency of the CIN on DOCSIS channels. This measurement is important because the transition network

has the potential to affect the latency budgets already established for DOCSIS devices. The latency measurement function is described in Section 8.6.1.

The DLM Sublayer Header is also used by the RPD to alert the CCAP Core that the RPD's output queue depth for a particular channel has reached or exceeded programmed threshold. The buffer depth monitoring alert function is described in Section 8.8.2 and in Section 8.6.2.

8.6.1 DEPI Latency Measurement

To perform the latency measurement, a packet is sent using the DEPI Latency Measurement Sublayer Header with Function '0000' (shown in Figure 59). The receiver responds to DLM requests. This sublayer header is designed to be used with any active L2TPv3 session between the CCAP Core and the RPD. It may be used with MPT, MCM, and PSP pseudowire types. It is anticipated that this particular message exchange may occur between hardware mechanisms at either end of the DEPI interface.




Figure 59 - DEPI Latency Measurement (DLM) Sublayer Header for Latency Measurements

Table 33 - DEPI Latency Measurement (DLM) Sublayer Header for Latency Measurement Details

Field	Description
V	1 bit. VCCV bit. Set to 0. Reserved for compatibility with [RFC 5085].
S	1 bit. Sequence bit. Set to 0.
H	2 bits. Extended Header bits. Set to '01' to indicate the presence of the DLM sublayer header.
X	1 bit. Reserved bit. Set to all zeroes at the transmitter, ignored at the receiver.
Flow ID	3 bits. Flow Identifier.
Reserved	4 bits. Reserved field. Set to all zeroes at the transmitter, ignored at the receiver.
Function	4 bits. Set to '0000' to indicate DEPI latency measurement function.
Code Field	1 byte. The permitted values are described below: 0 = A DLM-EI-RQ (DLM EQAM Ingress Request) packet originated by the CCAP Core requesting a measurement to be made at a reference point adjacent to the DEPI ingress port of the RPD. 1 = A DLM-EI-RP (DLM EQAM Ingress Reply) packet originated by the RPD with a Timestamp End value calculated at a reference point adjacent to the DEPI ingress port of the RPD. 2 = A DLM-EE-RQ (DLM EQAM Egress Request) packet originated by the CCAP Core requesting a measurement to be made at a reference point adjacent to the DEPI egress port of the RPD. 3 = A DLM-EE-RP (DLM EQAM Egress Reply) packet originated by the RPD with a Timestamp End value calculated at a reference point adjacent to the DEPI egress port of the RPD. Values 4-255 are reserved.
Transaction ID	1 byte. This is a unique ID assigned by the sender and returned by receiver. The transaction ID is unique for each transaction.
DOCSIS Timestamp Start	4 bytes. Timestamp sent by sender.
DOCSIS Timestamp End	4 bytes. Timestamp existing at the receiver.

The CCAP Core MUST support the sending of a DLM-EI-RQ packet and the receiving of a DLM-EI-RP packet on DEPI flows corresponding to DOCSIS channels. The CCAP Core MAY support the sending of a DLM-EE-RQ

packet and the receiving of a DLM-EE-RP packet on DEPI flows corresponding to DOCSIS channels. The CCAP Core MUST send the DLM packet to the RPD on an existing R-DEPI session flow (either PSP, MCM, or D-MPT). The CCAP Core MUST set the proper DSCP values for the DLM packet based on the R-DEPI session being measured.

The CCAP Core MUST NOT send a DLM-EI-RQ or DLM-EE-RQ packet to a particular DEPI session, if there is already a DLM-EI-RQ or DLM-EE-RQ outstanding for that session, or if a DOCSIS frame on that flow has been segmented and the complete DOCSIS frame has not been sent. The CCAP Core MUST place the current [DTI] 32-bit timestamp value in the DOCSIS Timestamp Start field of the message. The CCAP Core SHOULD use a timestamp value that is accurate to within 100 μs of the current timestamp as derived from [R-DTI].

The RPD MUST support the receiving of a DLM-EI-RQ packet and the sending of a DLM-EI-RP packet. The RPD MAY support the receiving of a DLM-EE-RQ packet and the sending of a DLM-EE-RP packet. The RPD MUST support the use of the DLM packet within any active DEPI session corresponding to a DOCSIS channel. The RPD is not required to support more than one concurrent latency measurement per session. The RPD MUST ensure that the timestamp value inserted in the DOCSIS Timestamp End field is accurate to within 100 μs of the current timestamp used for SYNC insertion and/or correction or Timestamp Message Block insertion on OFDM channels.

For DLM-EI-RQ packets, the RPD MUST perform the timestamp insertion prior to queuing the DEPI frame on the D-MPT or PSP QoS queues. For DLM-EE-RQ packets, if supported, the RPD MUST perform the timestamp insertion at the point where the SYNC message is originated; this is discussed in Section 6.1.3.2. The RPD MUST send the completed Latency Measurement Packet back to the CCAP Core that originated the measurement request, and do so within the RPD DLM Timeout value specified in Parameters and Constants.

An RPD that does not support a DLM-EE-RQ packet MUST silently discard the DLM packet without generating a DLM-EE-RP response packet.

For pseudowires carrying data for multiple channels, DLM-EI-RQ packets are intended to measure latency on a network path that's common for all channels of the pseudowire. Thus, the results of such measurements can be considered valid for the entire pseudowire or any of its channels.

The details of the operation of DLM-EE-RQ on pseudowires carrying data for multiple channels is outside of the scope of this specification.

8.6.2 DEPI Buffer Depth Alert Message

Section 8.8.2 describes buffer depth monitoring function implemented in the RPD. When the CCAP Core enables monitoring for a selected channel and the computed buffer depth average for this channel exceeds programmed threshold, the RPD sends a single L2TPv3 alert message as well as the GCP Notify message to the CCAP Core. The format of the L2TPv3 Buffer Depth Alert message is shown on Figure 60. This message uses DLM Sub-Layer Header with Function field set to '0001". The detailed description of the fields of this message is provided in Table 34.




Figure 60 - DLM Sub-layer Header for Buffer Depth Monitor Alert Function

Table 34 - Details of DLM Sublayer Header for Buffer Depth Monitor Alert Function

Field	Description
V	1 bit. VCCV bit. Set to 0. Reserved for compatibility with [RFC 5085].
S	1 bit. Sequence bit. Set to 0.

Field	Description
H	2 bits. Extended Header bits. Set to '01' to indicate the presence of the DLM sublayer header.
X	1 bit. Reserved bit. Set to all zeroes at the transmitter, ignored at the receiver.
Flow ID	3 bits. Set to zeroes at the transmitter, ignored at the receiver.
Reserved	4 bits. Reserved field. Set to all zeroes at the transmitter, ignored at the receiver.
Function	4 bits. Set to '0001' to indicate DEPI Buffer Depth Alert Message Format.
Reserved 2	1 byte. Reserved field. Set to all zeroes at the transmitter, ignored at the receiver.
Transaction ID	1 byte. This field is not used in Buffer Depth Monitor Message. This field is set to all zeroes at the transmitter, ignored at the receiver.
Channel ID/ MPTS_TAG	1 byte. Channel ID from the PSP segment header (or MPTS_TAG from the DEPI-MCM header) as communicated in Remote End ID AVP. For pseudowires carrying data for multiple channels this field identifies the specific channel. For pseudowires carrying data for a single channel this field is transmitted as all zeros.
Reserved 3	3 bytes. Reserved field. Set to all zeroes at the transmitter, ignored at the receiver.

8.7 DEPI Multicast Packet Header Format

The packet headers for the DEPI multicast packet are identical to the headers shown for the DEPI unicast packet in Figure 51. The only difference is that the source and destination IP address represent the Source and Group address of the multicast stream, and the Ethernet Destination address is also multicast instead of unicast.

8.8 CCAP Core Output Rate

It is possible that if the CCAP Core was to send packets at a data rate which exactly equaled the payload rate of the downstream QAM and those packets were subject to jitter, the RPD would insert an MPEG-TS null byte. Once the packets from the CCAP Core arrived, there would be an output queue delay that would equal the worst case jitter present at the DEPI input of the RPD. That queuing delay would not be removed until the input to the RPD was interrupted and the internal queue was drained. Furthermore, the rate at which this delay is removed is related to the amount of jitter the input stream has, the peak input rate, and the maximum burst size of the input stream.

Unlike SC-QAM channels, the downstream OFDM channel defined by DOCSIS 3.1 does not have a single measureable output rate. The actual output rate of a channel depends on a number of factors such as the proportion of the traffic sent to each of the defined profiles and the effectiveness of the PMD encoding. When more data is sent on the downstream profiles with higher effective modulation rates, the channel output rate increases. When more data is sent on downstream profiles with lower effective modulation rates, the channel output rate decreases.

The effective rate of a channel increases when the codeword encoder is able to assemble more full-length FEC codewords. While the CCAP Core is able to estimate the channel output rate based on the current traffic mix, it has no visibility into the PMD encoding process in the RPD device.

To prevent persistent buildup of buffers in the RPD, the CCAP Core MUST limit the output rate of the aggregate payload of all DEPI flows that are destined to the same DOCSIS downstream channel within an RPD. The CCAP Core MUST support OSSI-configurable de-rating of the aggregate DEPI payload to DOCSIS downstream channels with a default of 99% and a range of 90-99% with a granularity of one percent.

For all DOCSIS DEPI payload types, the CCAP Core MUST limit its burst of payload output at the CIN interface to a maximum of 1.5 msec of the estimated RF payload output rate. The value of 1.5 msec was chosen to exceed the RF transmission time of three (3) 1522-byte Packet PDUs at the slowest SC-QAM RF modulation rate of QAM-64. Any such burst is absorbed in the RPD jitter buffer as specified in Section 8.8.1.

For video and NDF DEPI payloads, the CCAP Core MUST maintain output with a deviation of no more than 1.0 msec ahead or behind the RF output rate of the payload.

The CCAP Core MAY provide vendor-specific configuration to increase the maximum DEPI payload burst.

The CCAP Core MUST be able to rate limit the aggregate of all DEPI flows, including any null MPEG packets that the CCAP Core may have inserted, that are destined to the same downstream channel within an RPD. The peak rate of this aggregate is configurable to be a percentage of the nominal or estimated downstream channel payload rate. The burst size of this aggregate is configurable (see [R-OSSI]). The default burst size of the aggregate is the equivalent of three frames per DEPI session. (For a frame size of 1522 bytes, this is 4566 bytes.)

8.8.1 Downstream Output Buffer History

The RPD samples the aggregate depth of the enqueued data waiting for transmission in its downstream channels' output buffers every 1 msec and stores the measurements in a circular buffer containing one thousand samples, thus providing buffer occupancy history for the most recent one second interval. In cases where the RPD maintains multiple output queues per channel, the RPD records the sum of the measured individual queue depths for all output queues for a particular channel. This sum does not include the data enqueued for transmission on PLC sub-channel. The CCAP Core can read the content of the measurement history buffer via GCP at any time. Based on this information, the CCAP Core can analyze such issues as latency and jitter, and optimize the channel output rate.

While the RPD measures the actual buffer occupancy in bytes, to minimize internal storage requirements and the volume of data exchanged via GCP, the buffer occupancy measurement samples are normalized to 8-bit values, where 0 represents an empty buffer and 255 represents a full buffer. The values of the buffer occupancy samples are the result of the division of the current length of the output queue measured in bytes by the normalization factor (NF). The CCAP Core provides an NF for each channel computed from the channel output rate and the output buffer depth. The method for calculation of the normalization factor is explained in the following examples.

Example 1:

For SC-QAM channels with an output rate of 38.8 Mbps with buffer depth equivalent to 8 msec of traffic at maximum rate, the maximum buffer depth, as measured in bytes can be calculated as

$$D = (8 \text{ [msec]} * 38.8 \text{ [Mbps]}) / 8 \text{ [B/b]} = 38.8 \text{ [kB]}$$

$$\text{Normalization Factor: } NF = D / 256 = 38.8 \text{ [kB]} / 256 = 151 \text{ [B]}$$

Example 2:

For an OFDM channel with a maximum output of 1 Gbps with buffer depth equivalent to 8 msec of traffic at maximum rate, the buffer depth can be calculated as

$$D = (8 \text{ [msec]} * 1 \text{ [Gbps]}) / 8 \text{ [B/b]} = 1 \text{ [MB]}$$

$$\text{Normalization Factor: } NF = D / 256 = 1 \text{ [MB]} / 256 = 3.9 \text{ [kB]}$$

Example 3:

For an MPEG video channel with an output rate of 38.8 Mbps with buffer depth equivalent to 10 msec of traffic at maximum rate, the maximum buffer depth, as measured in bytes can be calculated as

$$D = (10 \text{ [msec]} * 38.8 \text{ [Mbps]}) / 8 \text{ [B/b]} = 48.5 \text{ [kB]}$$

$$\text{Normalization Factor: } NF = D / 256 = 48.5 \text{ [kB]} / 256 = 189 \text{ [B]}$$

The RPD MUST support Output Buffer Occupancy history for downstream DOCSIS channels, MPEG video channels, and NDF channels.

8.8.2 Downstream Buffer Depth Monitoring Alerts

The CCAP can instruct the RPD to actively monitor the depth of its downstream queues and alert the CCAP Core via GCP if the average buffer depth for a particular channel exceeds a programmed threshold. This mechanism is designed to ensure that the buffer buildup issues within the RPD can be discovered in a timely manner so that the CCAP Core can correct them before buffers overflow and channel data is lost.

The CCAP enables Downstream Buffer Alerts by sending the RPD a Buffer Alert Threshold (BAT) value, a smoothing factor exponent "N" for a selected channel and by enabling monitoring via GCP. Upon enablement, the RPD starts computing the exponential moving average (EMA) of the aggregate output queue depth for the channel, sampling the queue depth every millisecond. When the computed value of the queue depth EMA exceeds the BAT value, the RPD notifies the CCAP Core a single time that the threshold has been exceeded. The notification is sent via L2TPv3 message (described in Section 8.6.2) as well as via GCP Notify message.

Upon reception of an alert, the CCAP Core takes a corrective action; it reduces the channel maximum output rate by a fraction. The CCAP Core determines the size of the rate reduction step. This specification recommends that the reduction step size to be at minimum one half of one percent of the current channel rate.

After the CCAP Core takes a corrective action, it can re-enable buffer depth monitoring alerts in the RPD again. Such an approach, which requires the CCAP Core to consecutively rearm the alert function after reception of an alert, is intended to ensure that the RPD does not flood the CCAP Core with a large number of alerts.

The RPD calculates the buffer depth EMA using the following equation

$$A_t = \alpha * D_t + (1 - \alpha) A_{t-1}$$

Where

A_t - current average, A_{t-1} - previous sample average

D - a normalized buffer depth (see Section 8.8.1 for the description of the normalization)

α - a smoothing factor, $\alpha = 1/2^N$, N is a parameter configured by the CCAP Core via GCP

The RPD SHOULD support Downstream Buffer Depth Monitoring Alerts for OFDM channels.

The RPD MAY support Downstream Buffer Depth Monitoring Alerts for SC-QAM DOCSIS channels.

The RPD MAY support Downstream Buffer Depth Monitoring Alerts for NDF channels.

The CCAP Core SHOULD support Downstream Buffer Depth Monitoring Alerts for OFDM channels.

The CCAP Core MAY support Downstream Buffer Depth Monitoring Alerts for SC-QAM DOCSIS channels.

The CCAP Core MAY support Downstream Buffer Depth Monitoring Alerts for NDR channels.

8.9 Bidirectional Forwarding Detection

Bidirectional Forwarding Detection (BFD) is the widely adopted IETF standard (see [RFC 5580]) that allows the two endpoints to monitor the reachability over a link or a network by both ends periodically sending "hello" packets to each other. It is an efficient and generic hello/keep-alive protocol that is used by many applications, and plays a key role in fast convergence and traffic redirection to achieve necessary SLA requirements of applications.

[RFC 5085] describes Connectivity Verification (CV) types using Bidirectional Forwarding Detection (BFD) with Virtual Circuit Connectivity Verification (VCCV). VCCV provides a control channel that is associated with an L2TPv3 pseudowire (PW), as well as the corresponding operations and management functions such as connectivity verification to be used over that control channel. Embedding VCCV information within the pseudowire permits the CV function to be handled by the same hardware components that are involved in processing the tunneled data.

Recently, the IETF focused on several aspects of BFD in order to further improve scalability and efficiency, to expand failure detection coverage and to allow BFD usage for more scenarios. These efforts resulted in creation of Seamless BFD (S-BFD) described in [S-BFD-BASE] and related IETF RFCs.

The RPD MUST support S-BFD responder functionality as specified in [S-BFD-BASE] and related L2TPv3 extensions specified in [S-BFD-VCCV] and [S-BFD-DSR-L2TP].

The RPD MUST support the ability to respond to S-BFD requests with a minimum rate of 100 responses per second under normal operational conditions.

Figure 61 shows the format of a BFD packet with a UDP header and the BFD Sublayer header.




Figure 61 - BFD Sublayer (with UDP)

8.10 IPv6 Support

The MHA_{v2} [DEPI] protocol supports both IPv4 and IPv6 network layer addressing. The RPD **MUST** support DEPI transport with IPv4 and IPv6 headers. The CCAP Core **MUST** support DEPI transport with IPv4 and IPv6 headers.

Additional considerations for processing of IPv6 extension headers are critical for IPv6 standard compliance and for supporting future capabilities of the RPD. Segment routing is one of the technologies considered for addressing QoS challenges in future SDN-managed access networks. The following requirements have been crafted in anticipation of deployment of segment routing over IPv6 in the CIN environment.

The Segment Routing impact on the RPD is fairly limited. The RPD does not need to understand the semantics of Source Routing. The RPD appends IPv6 extension headers to transmitted DEPI data packets or removes or bypasses those headers in received DEPI data packets.

In the ingress direction, the RPD's DEPI receiver **MUST** be able to process received packets with IPv6 extension headers up to 176 bytes long. The extension header length of 176 bytes was calculated based on the assumption that the longest Segment Routing extension header will contain a segment routing chain of 6 segments.

This specification does not define requirements for the interpretation or processing of the content of the extension headers with the exception for identification of the Hop-by-Hop header. The RPD's DEPI receiver **MUST** "punt" all packets with "Hop-by-Hop" extension headers to the RPD software for further processing.

The DEPI receiver in the RPD **MUST** be able to locate the L2TPv3 header in a packet based on traversing of the IPv6 extension header chain, as the RPD may have no prior knowledge of header size in the received packets.

In the egress direction, the DEPI transmitter in the RPD **MUST** be able to insert IPv6 extension headers up to 176 bytes long into transmitted DEPI packets. The number of unique IPv6 extension headers per RPD is limited to 4 per each CCAP Core the RPD communicates with.

The RPD **SHOULD** be capable to update the inserted extension headers (or the entire packet headers) atomically, that is without disrupting the flow of data packets. The RPD **MUST** minimize the disruption to the flow of transmitted data packets when the IPv6 headers are updated.

The method by which the RPD receives IPv6 headers to be inserted into transmitted DEPI packets will be defined in future revisions of R-PHY specifications.

Annex A Parameters and Constants (Normative)

Table 35 - Parameters and Constants

System	Name	Time Reference	Minimum Value	Default Value	Maximum Value
CCAP Core, RPD	SYNC Interval	Nominal time between transmission of SYNC messages.			200 msec
CCAP Core, RPD	HELLO Timer	L2TPv3 section 4.4		60 sec	
CCAP Core, RPD	Control Message Timeout	L2TPv3 section 4.2	First attempt 1 sec, then second attempt 2 sec, then third attempt 4 sec, then each additional attempt 8 sec.		
CCAP Core, RPD	Control Message Retry Count	L2TPv3 section 4.2			10
CCAP Core, RPD	StopCCN Timeout	L2TPv3 section 3.3.2	31 sec		
RPD	EQAM DLM Timer	Time between receipt and retransmission of a DLM packet.			100 msec

Annex B DOCS-IF-M-CMTS-MIB (Normative)

OSS updates will be incorporated into the [R-OSSI] specification.

Annex C Format and Content for Event, SYSLOG, and SNMP Notification (Normative)

Updates will likely be incorporated into the [R-OSSI] specification.

Annex D DEPI Path Redundancy (Normative)

The MHA_{v2} updates to DEPI Path Redundancy will be addressed in a future revision.

Appendix I DEPI and DOCSIS System Performance (Informative)

This appendix will be updated in a future revision.

Appendix II Acknowledgements

On behalf of the cable industry and our member companies, CableLabs would like to thank the following individuals for their contributions to the development of this specification.

Contributor	Company Affiliation
John T. Chapman, Sangeeta Ramakrishnan, Pawel Sowinski, Alon Bernstein, Brian Bresnahan, Carlos Pignataro, Michael Patrick	Cisco Harmonic

On behalf of the cable industry and our member companies, CableLabs would like to thank the following individuals for their contributions to the development of the technology and participation in the Remote PHY Working Group.

Contributor	Company Affiliation
Bill Powell	Alcatel-Lucent
Brian Kurtz	Altera
Carlton Lane, Linda Mazaheri	Analog
Tom Ferreira, Steve Foley, Anand Goenka, Jeff Howe, Hari Nair	Arris
Andrew Chagnon, Victor Hou, Niki Pantelias, David Pullen	Broadcom
Stuart Hoggan, Volker Leisse, Jon Schnoor, Karthik Sundaresan, Nikhil Tayal, Jun Tian	CableLabs
Andrew Sundelin	CableLabs Consultant
Naor Goldman	Capacicom
Dave Fox, Maike Geng	Casa Systems
David Claussen	Charter
Nobo Akiya, Alon Bernstein, Brian Bresnahan, John T. Chapman, Hang Jin, Tong Liu, Carlos Pignataro, Sangeeta Ramakrishnan, John Ritchie, Pawel Sowinski, Don Strausberger, Yi Tang, Bill Wall, Gerry White	Cisco
Philippe Perron	Cogeco
John Bevilacqua, Nagesh Nandiraju, Saifur Rahman, Jorge Salinger, Joe Solomon, Douglas Will	Comcast
Jeff Ford, Al Garrett	Complex IQ
Ony Anglade, Mike Cooper	Cox Communications
Samir Parikh	Gainspeed Networks
João Campos, Even Kristoffersen	Get
Adi Bonen, Michael Patrick	Harmonic
Jim Chen, Hesham ElBakoury, Karl Moerder, Jack Moran, Guangsheng Wu	Huawei
Phil Oakley	LGI
Stan Bochenek, Ajay Kuckreja	Maxim Integrated
Len Dauphinee, David Huang, Louis Park, Sridhar Ramesh, Patrick Tierney, Scott Walley	MaxLinear
Rei Brockett	Pace/Aurora
Nasir Ansari, George Hart	Rogers
Kevin Kwasny	Shaw
Lee Johnson	ST Micro
Paul Brooks, Kirk Erichsen	Time Warner Cable
Colin Howlett, Douglas Johnson	Vecima
Faten Hijazi, Alex Luccisano	Xilinx

Karthik Sundaresan and Jon Schnoor, CableLabs

Appendix III Revision History

Engineering Changes for CM-SP-R-DEPI-I02-151001

ECN	Date Accepted	Summary	Author
R-DEPI-N-15.1358-2	9/9/2015	Clarifications for PLC Message Block transport over DEPI	Pantelias
R-DEPI-N-15.1362-3	9/9/2015	DEPI Signaling for multichannel pseudowires	Sowinski
R-DEPI-N-15.1363-2	9/9/2015	Video Synchronization Mode	Ramakrishnan

Engineering Change for CM-SP-R-DEPI-I03-160121

ECN	Date Accepted	Summary	Author
R-DEPI-N-15.1404-2	12/9/2015	R-PHY pseudowire subtypes	Sowinski

Engineering Changes for CM-SP-R-DEPI-I04-160512

ECN	Date Accepted	Summary	Author
R-DEPI-N-16.1462-1	4/21/2016	Number of DEPI multicast addresses	Sundaresan
R-DEPI-N-16.1481-1	4/21/2016	R-DEPI Omnibus	Sundaresan

Engineering Changes for CM-SP-R-DEPI-I05-160923

ECN	Date Accepted	Summary	Author
R-DEPI-N-16.1550-1	8/4/2016	DEPI Multicast Support (replaces R-DEPI-16.1532)	Pantelias
R-DEPI-N-16.1574-1	9/1/2016	R-DEPI Null MPEG Frames for lost D-MPT Packets	Bonen
R-DEPI-N-16.1577-1	9/1/2016	Output buffer monitoring and buffer depth alerts	Sowinski
R-DEPI-N-16.1580-3	9/1/2016	DEPI miscellaneous technical changes.	Sowinski
R-DEPI-N-16.1581-1	9/1/2016	DEPI miscellaneous changes	Sowinski

Engineering Changes for CM-SP-R-DEPI-I06-170111

ECN	Date Accepted	Summary	Author
R-DEPI-N-16.1632-3	11/23/2016	DEPI Control Plane Modification for PHB and UDP Port signaling	Sowinski
R-DEPI-N-16.1650-1	11/23/2016	R-DEPI - move technical requirements out of table cells	Schnoor
R-DEPI-R-16.1672-2	12/15/2016	Cleanup and clarifications of several requirements	Sowinski

Engineering Changes for CM-SP-R-DEPI-I07-170524

ECN	Date Accepted	Summary	Author
R-DEPI-N-17.1695-2	3/9/2017	Discrepancies of specs in R-PHY-I06 versus R-DEPI-I06 for MCM and DMPT for DOCSIS (RPHY-161) part2	Huang
R-DEPI-N-17.1698-1	3/16/2017	Clarify device requirements for support of UDP headers.	Sowinski

ECN	Date Accepted	Summary	Author
R-DEPI-N-17.1701-1	3/23/2017	Remove references to R-DTI Client Test Port	Armstrong
R-DEPI-N-17.1705-1	3/30/2017	Add MUST requirements under Table 7-5 for the "required" line items in the table.	Schnoor
R-DEPI-N-17.1706-1	3/30/2017	Fix table notes in DEPI spec	Sowinski
R-DEPI-N-17.1723-2	4/13/2017	R-DEPI Editorial Issues	Sowinski