

Data-Over-Cable Service Interface Specifications Business Services over DOCSIS®

Layer 2 Virtual Private Networks

CM-SP-L2VPN-I14-150305

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.

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-L2VPN-I14-150305			
Document Title:	Layer 2 Virtual Private Networks			
Revision History:	I01 – Released March 28, 2006 I02 – Released July 28, 2006 I03 – Released December 22, 2006 I04 – Released May 18, 2007 I05 – Released August 3, 2007 I06 – Released December 6, 2007 I07 – Released February 15, 2008 I08 – Released May 22, 2008 I09 – Released June 11, 2010 I10 – Released October 04, 2012 I11 – Released August 8, 2013 I12 – Released November 20, 2013 I13 – Released April 3, 2014 I14 – Released March 5, 2015			
Date:	March 5, 2015			
Status:	Work in Progress	Draft	Issued	Closed
Distribution Restrictions:	Focus Team Only	CL/Member	CL/Member/Vendor	Public

Key to Document Status Codes

Work in Progress	An incomplete document, designed to guide discussion and generate feedback that may include several alternative requirements for consideration.
Draft	A document in specification format considered largely complete, but lacking review by Members and vendors. 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 for Certification testing if applicable. Issued Specifications are subject to the Engineering Change Process.
Closed	A static document, reviewed, tested, validated, and closed to further engineering change requests to the specification through CableLabs.

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.

Contents

1	INTRODUCTION	9
1.1	Requirements	9
1.2	Conformance	9
1.3	Conventions	10
2	REFERENCES	11
2.1	Normative References	11
2.2	Informative References	12
2.3	Reference Acquisition	13
3	TERMS AND DEFINITIONS	14
4	ABBREVIATIONS AND ACRONYMS.....	16
5	THEORY OF OPERATION (INFORMATIVE).....	21
5.1	L2VPN Features	21
5.1.1	Transparent LAN Service	21
5.1.2	Multiple ISP L2VPNs.....	23
5.1.3	Management L2VPNs	23
5.1.4	Other L2VPN-enabled Features	24
5.1.5	High Availability	24
5.2	CMTS Layer 2 Forwarding Architecture.....	25
5.2.1	L2VPN and Non-L2VPN Forwarding.....	25
5.2.2	Point-to-Point and Multipoint L2VPN Forwarding Modes	26
6	L2VPN OPERATION (NORMATIVE)	29
6.1	CMTS Bridging Model Requirements	29
6.2	Configuring L2VPN Forwarding	29
6.2.1	VPNID Subtype	32
6.2.2	Downstream Classifier L2VPN Encoding	32
6.2.3	L2VPN SA-Descriptor Subtype	33
6.2.4	Vendor-Specific L2VPN Encoding.....	34
6.2.5	Configuration Error Requirements.....	34
6.2.6	Network System Interface (NSI) Encapsulation.....	34
6.2.7	Virtual Private LAN Service (VPLS) and Virtual Private Wire Service	36
6.3	CMTS Upstream L2VPN Forwarding	40
6.4	CMTS Downstream L2VPN Forwarding	41
6.4.1	Multipoint Downstream Forwarding.....	43
6.4.2	DOCSIS 3.0 L2VPN Downstream Multicast Forwarding	43
6.5	L2VPN Isolation and Privacy	43
6.5.1	Protecting L2VPN Traffic.....	43
6.5.2	Preventing Leaking of non-L2VPN Traffic	44
6.6	CM and eSAFE Exclusion.....	45
6.6.1	CM and eSAFE Host Forwarding Model	46
6.6.2	Cable Modem MAC Bridge Interface Masks	47
6.6.3	Embedded Host Exclusion	47
6.6.4	CMTS embedded host MAC Address Learning.....	48
6.6.5	Interface-based Classification	48
6.7	L2VPN Quality Of Service	49
6.7.1	DOCSIS QoS.....	49
6.7.2	Backbone Network QoS.....	49
6.8	Stacked 802.1Q Tags or Tag-in-Tag operation	51
6.9	Spanning Tree and Loop Detection	52
6.10	High Availability	52

6.10.1	802 Encapsulation - Active/Standby Layer 2 Trunk Ports	52
6.10.2	802 Encapsulation - Link Aggregation	53
6.10.3	MPLS Encapsulation High Availability	53
6.11	MPLS Encapsulation - PW Redundancy	53
7	CABLE MODEM REQUIREMENTS	54
8	SERVICE OPERATIONS, ADMINISTRATION, AND MAINTENANCE (OAM)	56
8.1	Introduction	56
8.2	Service OAM Configuration	57
8.2.1	CM	57
8.2.2	CMTS	58
8.3	Fault Management	58
8.3.1	Continuity Check Messages (CCM)	58
8.3.2	Loopback	60
8.3.3	Linktrace	61
8.4	Performance Management	62
8.4.1	Frame Delay	62
8.4.2	Synthetic Frame Loss	63
9	LAYER 2 CONTROL PROTOCOL HANDLING	65
ANNEX A	CMTS DOCS-L2VPN-MIB REQUIREMENTS (NORMATIVE)	66
A.1	DOCS-L2VPN-MIB Conformance	66
A.2	DOCS-L2VPN-MIB Definitions	70
ANNEX B	PARAMETER ENCODINGS (NORMATIVE)	92
B.1	Capabilities	92
B.1.1	L2VPN Capability	92
B.1.2	Embedded Service/Application Functional Entity (eSAFE) Host Capability	92
B.1.3	Downstream Unencrypted Traffic (DUT) Filtering	92
B.2	Downstream Unencrypted Traffic (DUT) Filtering Encoding	92
B.2.1	Downstream Unencrypted Traffic (DUT) Control	93
B.2.2	Downstream Unencrypted Traffic (DUT) CMIM	93
B.3	L2VPN Encoding	93
B.3.1	VPN Identifier	94
B.3.2	NSI Encapsulation Subtype	95
B.3.3	eSAFE DHCP Snooping	100
B.3.4	CM Interface Mask (CMIM) Subtype	100
B.3.5	Attachment Group ID	101
B.3.6	Source Attachment Individual ID	101
B.3.7	Target Attachment Individual ID	101
B.3.8	Upstream User Priority	102
B.3.9	Downstream User Priority Range	102
B.3.10	L2VPN SA-Descriptor Subtype	102
B.3.11	Vendor Specific L2VPN Subtype	103
B.3.12	Pseudowire ID (deprecated)	103
B.3.13	Pseudowire Type	103
B.3.14	L2VPN Mode	103
B.3.15	Tag Protocol Identifier (TPID) Translation	104
B.3.16	L2CP Processing	105
B.3.17	DAC Disable/Enable Configuration	106
B.3.18	Pseudowire Class	106
B.3.19	Service Delimiter	106
B.3.20	Virtual Switch Instance Encoding	108
B.3.21	BGP Attribute sub TLV	109
B.3.22	VPN-SG Attribute sub TLV	110

B.3.23	<i>Pseudowire Signaling</i>	110
B.3.24	<i>L2VPN SOAM Subtype</i>	110
B.3.25	<i>Network Timing Profile Reference</i>	114
B.3.26	<i>L2VPN DSID</i>	114
B.4	Confirmation Codes	115
B.5	L2VPN Error Encoding	115
B.5.1	<i>L2VPN Errored Parameter</i>	115
B.5.2	<i>L2VPN Error Code</i>	116
B.5.3	<i>L2VPN Error Message</i>	116
B.6	CM Interface Mask Classification Criteria	116
B.7	L2VPN MAC Aging Encoding	117
B.7.1	<i>L2VPN MAC Aging Mode</i>	118
APPENDIX I EXAMPLE L2VPN ENCODINGS (INFORMATIVE)		119
I.1	Point-to-Point Example	119
I.2	Multipoint Example	122
I.3	Upstream L2VPN Classifier Example	125
APPENDIX II IEEE 802.1Q ENCAPSULATION (INFORMATIVE)		127
APPENDIX III EMBEDDED VLAN CM BRIDGING MODEL (INFORMATIVE)		128
III.1	IEEE 802.1Q and Embedded VLAN Model	129
III.2	Embedded Bridge MAC Domain Service Primitives	130
APPENDIX IV L2VPN NON-COMPLIANT CM RESTRICTIONS (INFORMATIVE)		132
IV.1	Leaking through non-compliant CMs	132
APPENDIX V ACKNOWLEDGEMENTS (INFORMATIVE)		133
APPENDIX VI REVISION HISTORY (INFORMATIVE)		134
VI.1	Engineering Change for CM-SP-L2VPN-I02-060728	134
VI.2	Engineering Change for CM-SP-L2VPN-I03-061222	134
VI.3	Engineering Change for CM-SP-L2VPN-I04-070518	134
VI.4	Engineering Changes for CM-SP-L2VPN-I05-070803	134
VI.5	Engineering Changes for CM-SP-L2VPN-I06-071206	134
VI.6	Engineering Change for CM-SP-L2VPN-I07-080215	134
VI.7	Engineering Changes for CM-SP-L2VPN-I08-080522	135
VI.8	Engineering Changes for CM-SP-L2VPN-I09-100611	135
VI.9	Engineering Changes for CM-SP-L2VPN-I10-121004	135
VI.10	Engineering Change for CM-SP-L2VPN-I11-130808	135
VI.11	Engineering Change for CM-SP-L2VPN-I12-131120	135
VI.12	Engineering Change for CM-SP-L2VPN-I13-140403	135
VI.13	Engineering Change for CM-SP-L2VPN-I14-150305	136

Figures

Figure 5–1 - Transparent LAN Service	21
Figure 5–2 - TLS over a MPLS Network	23
Figure 5–3 - CMTS L2VPN and Non-L2VPN Forwarding	25
Figure 5–4 - Point-to-Point Forwarding Mode	27
Figure 5–5 - Multipoint L2VPN Forwarding Mode Example	28
Figure 6–1 - CM, eMTA, and CPE Forwarding	46
Figure 8–1 - Operator and EVC MEG Level.....	58
Figure 8–2 - CCM in the MSO Network	59
Figure 8–3 - Loopback in the MSO Network	60
Figure 8–4 - Linktrace functionality in the MSO Network	61
Figure 8–5 - Frame Delay functionality in the MSO Network	63
Figure 8–6 - Frame Loss Functionality in the MSO Network	64
Figure I–1 - Point-To-Point L2VPN Traffic Forwarding Example	119
Figure I–2 - Multipoint L2VPN Forwarding of Traffic Example.....	122
Figure II–1 - Ethernet 802.1Q Tags.....	127
Figure III–1 - L2VPN Embedded VLAN (eVLAN) Model	128

Tables

Table 6–1 - L2VPN Encoding Subtype Location Summary	30
Table 6–2 - Cable Modem MAC Bridge Interfaces	47
Table 6–3 - eSAFE DHCP Snooping Substrings	48
Table 8–1 - High Level Operations and References.....	56
Table 8–2 - Service OAM Terminology.....	57
Table 8–3 - Flags Format in CCM PDU.....	59
Table 9–1 - L2CP Frames and Actions.....	65
Table I–1 - Point-to-Point CM1 L2VPN Encoding	120
Table I–2 - Point-to-Point CM2 L2VPN Encoding	120
Table I–3 - Point-to-Point CM3 L2VPN Encoding	121
Table I–4 - Multipoint CM1 L2VPN Encoding	123
Table I–5 - Multipoint CM2 L2VPN Encoding	123
Table I–6 - Multipoint CM3 L2VPN Encoding	124
Table I–7 - Multipoint CM4 L2VPN Encoding	125
Table I–8 - Upstream L2VPN Classifier Encoding	125
Table IV–1 - Non-L2VPN leaking through Non-compliant CMs configured for L2VPN	132

This page intentionally left blank.

1 INTRODUCTION

This specification describes requirements on both CMTSs and CMs in order to implement a DOCSIS Layer-2 Virtual Private Network (DOCSIS L2VPN) feature.

The L2VPN feature allows cable operators to offer a Layer 2 Transparent LAN Service (TLS) to commercial enterprises, which is one of the principal goals of the CableLabs Business Services over DOCSIS (BSoD) initiative.

In order to speed time to market, [L2VPN-DG] offers guidelines to CMTS manufacturers as to how to phase the implementation of requirements defined in this specification. Phase designations are only applicable to CMTS products. Cable modems are expected to support all required L2VPN features in Phase 1. ¹

1.1 Requirements

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.

Some normative statements require a CM or CMTS to silently ignore a condition which may be defined in future specifications. A requirement to silently ignore a condition means that the CM or CMTS:

- MAY increment a vendor-specific statistic;
- MUST NOT generate a log message; and
- MUST otherwise ignore the condition and continue operation as if the condition did not occur.

1.2 Conformance

A DOCSIS CMTS that claims to implement the DOCSIS L2VPN feature MUST implement the normative provisions of this document. A DOCSIS CM that claims conformance for DOCSIS L2VPN feature MUST implement the normative requirements of this document.

A CMTS or CM implementing this specification is said to be L2VPN compliant. For the remainder of this document, all references to a CMTS refer to an L2VPN compliant CMTS. A CM that has not implemented this specification is termed an L2VPN non-compliant CM.

¹ Paragraph added per L2VPN-07.0592-1 on 1/24/08 by KN.

An L2VPN-compliant CMTS **MUST** support an L2VPN non-compliant CM. This permits a cable operator to offer L2VPN subscriber service with currently deployed non-compliant CMs. Use of non-compliant CMs involves certain limitations that are detailed in Appendix IV. Using compliant CMs for L2VPN subscriber service avoids those limitations. The requirements for CMs to comply with this L2VPN specification are summarized in Section 7.

1.3 Conventions²

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

² Added per L2VPN-N-12.1065-11 on 9/5/12 by JB.

2 REFERENCES

2.1 Normative References ³

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 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 non-specific reference, the latest version applies.

[802.1AX]	IEEE 802.1AX, IEEE Standard for Local and metropolitan area networks--Link Aggregation, 2008.
[802.1D]	IEEE 802.1D, IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges, June 2004.
[802.1Q]	IEEE Std 802.1Q, IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Local Area Networks, August 2011.
[BPI+]	Baseline Privacy Plus Interface Specification, CM-SP-BPI+-C01-081104, November 4, 2008, Cable Television Laboratories, Inc.
[DOCSIS MAC]	Refers to [MULPIv3.1], [MULPIv3.0] and [RFI 2.0].
[ID LDPv6]	IETF Internet draft, Updates to LDP for IPv6, draft-ietf-mpls-ldp-ipv6-15, January 11, 2015, http://tools.ietf.org/html/draft-ietf-mpls-ldp-ipv6-15
[MEF 30.1]	MEF Technical Specification 30.1: Service OAM Fault Management Implementation Agreement: Phase 2, April 2013.
[MEF 31]	MEF Technical Specification 31: Service OAM Fault Management Definition of Managed Objects, January 2011.
[MEF 35]	MEF Technical Specification 35: Service OAM Performance Monitoring Implementation Agreement, April 2012.
[MEF 36]	MEF Technical Specification 36: Service OAM SNMP MIB for Performance Monitoring, January 2012.
[MULPIv3.0]	DOCSIS MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0-I26-150305, March 5, 2015, Cable Television Laboratories, Inc.
[MULPIv3.1]	DOCSIS MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I04-141218, December 18, 2014, Cable Television Laboratories, Inc.
[OSSIV3.0]	Operations Support System Interface Specification, CM-SP-OSSIV3.0-I25-150305, March 5, 2015, Cable Television Laboratories, Inc.
[PWTYPES]	MPLS Pseudowire Types Registry, March 30, 2010, Internet Assigned Numbers Authority.
[RFC 2544]	IETF RFC 2544, Benchmarking Methodology for Network Interconnect Devices, March 1999.
[RFC 2918]	IETF RFC 2918, Route Refresh Capability for BGP-4, September 2000.
[RFC 3031]	IETF RFC3031, Multiprotocol Label Switching Architecture, January 2001.
[RFC 3032]	IETF RFC 3032, MPLS Label Stack Encoding, January 2001.

³ Revised per L2VPN-N-07.0460-1, L2VPN-N-07.0471-2 on 7/5/07 by KN and L2VPN-N-10.0918-2 on 5/10/10. Revised per L2VPN-N-12.1065-11 on 9/5/12, per L2VPN-N-14.1135-5 on 3/27/14 by JB, and L2VPN-N-14.1186-1 on 1/15/15 by PO.

[RFC 3931]	IETF RFC 3931, Layer Two Tunneling Protocol - Version 3 (L2TPv3), March 2005.
[RFC 4385]	IETF RFC 4385, Pseudowire Emulation Edge-to-Edge (PWE3). Control Word for Use over an MPLS PSN, February 2006.
[RFC 4447]	IETF RFC 4447, Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP), April 2006.
[RFC 4448]	IETF RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks, April 2006.
[RFC 4667]	IETF RFC 4667, Layer 2 Virtual Private Network (L2VPN) Extensions for Layer 2 Tunneling Protocol (L2TP), September 2006.
[RFC 4761]	IETF RFC 4761, Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling, January 2007.
[RFC 4762]	IETF RFC 4762, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling, January 2007.
[RFC 4893]	IETF RFC 4893, BGP Support for Four-octet AS Number Space, May 2007.
[RFC 5036]	IETF RFC 5036, LDP Specification, October 2007.
[RFC 5286]	IETF RFC 5286, Basic Specification for IP Fast Reroute: Loop-Free Alternates, September 2008.
[RFC 5925]	IETF RFC 5925, The TCP Authentication Option, June 2010.
[RFC 6074]	IETF RFC 6074, Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs), January 2011.
[RFC 6624]	IETF RFC 6624, Layer 2 Virtual Private Networks Using BGP for Auto-Discovery and Signaling, May 2012.
[RFI 2.0]	DOCSIS Radio Frequency Interface Specification, CM-SP-RFIV2.0-C02-090422, April 22, 2009, Cable Television Laboratories, Inc.
[Y.1564]	ITU-T Recommendation Y.1564 (03/11), Ethernet service activation test methodology.
[Y.1731]	ITU-T Recommendation Y.1731 (11/13), OAM functions and mechanisms for Ethernet based networks.

2.2 Informative References ⁴

This specification uses the following informative references.

[DPoE ARCHv1.0]	DPoE-SP-ARCHv1.0-I03-140807, August 7, 2014, DOCSIS Provisioning of EPON, DPoE Architecture Specification, Cable Television Laboratories, Inc.
[DPoE DEMARCV1.0]	DPoE-SP-DEMARCV1.0- I04-140807, August 7, 2014, DOCSIS Provisioning of EPON, DPoE Demarcation Device Specification, Cable Television Laboratories, Inc.
[DPoE IPNEv1.0]	DPoE-SP-IPNEv1.0- I08-140807, August 7, 2014, DOCSIS Provisioning of EPON, IP Network Element Requirements, Cable Television Laboratories, Inc.
[DPoE MEFv2.0]	DPoE-SP-MEFv2.0-I03-140807, August 7, 2014, DOCSIS Provisioning of EPON, Metro Ethernet Forum Specification, Cable Television Laboratories, Inc.

⁴ Revised per L2VPN-N-07.0461-2 on 7/5/07, L2VPN-N-07.0592-1 on 1/24/08 by KN and L2VPN-N-10.0918-2 on 5/10/10. Revised per L2VPN-N-12.1065-11 on 9/5/12 by JB, and L2VPN-N-14.1186-1 on 1/15/15 by PO.

[DPoE MULPIv1.0]	DPoE-SP-MULPIv1.0-I09-140807, August 7, 2014, DOCSIS Provisioning of EPON, DPoE MAC and Upper Layer Protocols Requirements, Cable Television Laboratories, Inc.
[eDOCSIS]	eDOCSIS Specification, CM-SP-eDOCSIS-I28-150305, March 5, 2015, Cable Television Laboratories, Inc.
[eRouter]	Data-Over-Cable Service Interface Specifications, eRouter Specification, CM-SP-eRouter-I14-150305, March 5, 2015, Cable Television Laboratories, Inc.
[L2VPN-DG]	L2VPN Development Guidelines Technical Report, CM-TR-L2VPN-DG-V02-121206, December 6, 2012, Cable Television Laboratories, Inc.
[MEF 6.1.1]	MEF 6.1.1, Layer 2 Control Protocol Handling Amendment to MEF 6.1, January 2012.
[PKT-PROV]	PacketCable Provisioning, PKT-SP-PROV-C01-071129, November 29, 2007, Cable Television Laboratories, Inc.
[RFC 2685]	IETF RFC 2685, Virtual Private Network Identifier, September 1999.
[RFC 3107]	IETF RFC 3107, Carrying Label Information in BGP-4, May 2001.
[RFC 3209]	IETF RFC 3209, RSVP-TE: Extensions to RSVP for LSP Tunnels, December 2001.
[RFC 3270]	IETF RFC 3270, Multi-Protocol Label Switching (MPLS) Support of Differentiated Services, May 2002.
[RFC 3985]	IETF RFC 3985, Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture, March 2005.
[RFC 4363]	IETF RFC 4363, Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions, January 2006.
[RFC 4364]	IETF RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs), February 2006.
[RFC 4664]	IETF RFC 4664, Framework for Layer 2 Virtual Private Networks (L2VPNs), September 2006.

2.3 Reference Acquisition ⁵

- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027; Phone +1-303-661-9100; Fax +1-303-661-9199; Internet: <http://www.cablelabs.com>
- Institute of Electrical and Electronics Engineers (IEEE), Internet: <http://www.ieee.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, Internet: <http://www.ietf.org>
- International Organization for Standardization (ISO), 1, rue de Varembe, Case postale 56, CH-1211 Geneva 20, Switzerland, Phone +41 22 749 01 11; Fax +41 22 733 34 30; Internet: <http://www.iso.org>
- International Telecommunication Union (ITU), Place des Nations, CH-1211, Geneva 20, Switzerland; Phone +41-22-730-51-11; Fax +41-22-733-7256. Internet: <http://www.itu.int>
- Internet Assigned Numbers Authority (IANA), Internet: <http://www.iana.org/>
- Metro Ethernet Forum (MEF), 6033 W. Century Blvd, Suite 830, Los Angeles, CA 90045 Phone +1-310-642-2800; Fax +1-310-642-2808. Internet: <http://metroethernetforum.org>

⁵ Revised per L2VPN-N-12.1065-11 on 9/5/12 by JB.

3 TERMS AND DEFINITIONS

This specification uses the following terms in addition to those defined in [RFI 2.0].⁶

Bridged Network	A set of IEEE 802 LANs interconnected by IEEE 802.1D MAC bridges.
Compliant CM	A CM that implements this DOCSIS L2VPN specification.
DOCSIS L2PDU	A Packet PDU of a DOCSIS MAC Frame, i.e., the L2PDU following a MAC Header with FC_TYPE=00. This definition means that a MAC Management message with FC_TYPE=11 is <i>not</i> considered to be a DOCSIS L2PDU, even though the form of a MAC Management Message Header is the same form as an L2PDU.
DOCSIS MAC Frame	The unit of transmission on the DOCSIS cable RF interface, consisting of a MAC Header and a (possibly null) Data PDU. The FC_TYPE field of MAC Header identifies the Data PDU as either a Packet PDU (FC_TYPE=00), or a MAC-specific PDU (FC_TYPE=11).
Downstream Service Identifier (DSID)	A 20-bit value in a DOCSIS extended header that identifies a stream of packets distributed to the same cable modem or group of cable modems. The DSID value is unique within a MAC Domain. For sequenced packets, the DSID identifies the resequencing context for downstream packet bonding in the CM.
Flooding	An operation of an L2 Bridge in which it replicates an L2PDU addressed to a group MAC or unlearned individual MAC address to all Bridge Ports other than the L2PDU's ingress port.
Group MAC (GMAC) Address	An IEEE 6-byte MAC address with the first transmitted bit (the group bit) set to 1, indicating that the address refers to a group of MAC hosts. In the canonical representation of MAC addresses used for Ethernet transmission, the Group bit is the least significant bit of the first byte. The all-1s broadcast MAC address is considered to be a GMAC address.
Individual MAC Address	An IEEE 6-byte MAC address with the first transmitted bit (the group bit) set to 0, indicating that the address refers to a single MAC host. For the Ethernet MAC addresses of DOCSIS, the group bit is the least significant bit of the first byte of the MAC address.
L2 Forwarder	A network element that forwards layer 2 packets from one L2 interface to another L2 interface. A Layer 2 Forwarder may operate in Point-to-Point or Multipoint forwarding mode, i.e., forwarding between only two interfaces without learning; or Multipoint, forwarding unicast-destined packets only to the interface from which a MAC address was learned.
L2 Interface	A physical interface port or virtual circuit on which an L2PDU is transmitted. Physical L2 interface ports include an Ethernet NSI at a CMTS or the CMCI port at a CM. Virtual circuit L2 Interfaces include a CMTS Network System Interface (NSI) PseudoWire (PW) and a CMTS single-CM BPI Security Association. An L2 Interface may or may not have an ifIndex assigned to it.
L2 Protocol Data Unit (L2PDU)	A sequence of bytes consisting of a Destination MAC Address (DMAC), Source MAC Address (SMAC), (optional) Tag Header(s), EtherType/Length, L2 Payload, and CRC.

⁶ Revised per L2VPN-N-07.0504-1 on 11/9/07 by KN, and L2VPN-N-14.1186-1 on 1/15/15 by PO.

L2 Virtual Private Network (L2VPN)	A set of LANs and the L2 Forwarders between them that enable hosts attached to the LANs to communicate with Layer 2 Protocol Data Units (L2PDUs). A single L2VPN forwards L2PDUs based only on the Destination MAC (DMAC) address of the L2PDU, transparent to any IP or other Layer 3 address. A cable operator administrative domain supports multiple L2VPNs, one for each subscriber enterprise to which Transparent LAN Service is offered.
L2VPN Identifier	An octet string that uniquely identifies an L2VPN within a cable operator administrative domain, corresponding to a single subscriber enterprise.
L3 Forwarder	A network element that forwards a Layer 3 PDU from an ingress interface to one or more egress interfaces. Also called a Router.
Learning	An operation of a layer 2 Bridge by which it associates the Source MAC (SMAC) address of an incoming L2PDU with the Bridge Port from which it arrived.
Management L2VPN	An L2VPN for the post-registration SNMP traffic to eCM or eSAFE devices. May be combined with a Provisioning L2VPN.
Multipoint L2 Forwarding	Operation of an L2 Forwarder among multiple L2 networks that forwards individual MAC destined packets only to the interface from which a source MAC address was learned and that floods group MAC destined packets to all interfaces.
Non-compliant CM	A CM that does not implement this DOCSIS L2VPN specification.
Point-to-Point L2 Forwarding	Operation of an L2 Forwarder between only two L2 networks with no source MAC address learning.
Provisioning L2VPN	An L2VPN for the pre-registration traffic of DHCP, TOD, and TFTP that provisions eCMs and eSAFE hosts. May be combined with a Management L2VPN.
Security Association (SA)	An association between the CMTS and a set of CMs in a MAC Domain that enables encrypted communication between the CMTS and the CM set. A Single CM SA is one with a single CM, and enables a private point-to-point L2 Network connection between the CMTS and the CPE LAN of that CM. A Security Association Descriptor (SA-Descriptor) is a multiple-part message element defined in the DOCSIS Baseline Privacy specification [BPI+] that includes a Security Association ID (SAID).
Security Association ID (SAID)	A 14-bit identifier that appears in a BPI Extended Header (BPI-EH) of a DOCSIS PDU packet to identify the key used to encrypt the packet.
Tag Header	A 16-bit Tag Protocol ID (0x8100) followed by a 16-bit Tag Control field. The Tag Control field consists of a 3-bit User Priority field, a 1-bit Canonical Format Indicator, and a 12-bit VLAN ID [802.1Q].
Transparent LAN Service (TLS)	A service offering of a cable operator that implements a private L2VPN among the CPE networks of the CMs of single subscriber enterprise.
Virtual LAN (VLAN)	A <i>subset</i> of the LANs of an IEEE 802.1 Bridged Network to which a VLAN Identifier (VLAN ID) is assigned. An L2VPN may consist of several VLANs, each with different VLAN IDs, and even of VLANs on different IEEE 802.1 Bridged Networks with the same VLAN ID.
Virtual LAN Identifier (VLAN ID)	An IEEE 802.1Q VLAN ID is a 12-bit number that identifies a VLAN within an IEEE 802.1 Bridged Network. An IEEE 802.1Q (formerly known as 802.1ah) stacked VLAN ID consists of an outer Service 12-bit VLAN ID and an inner Customer 12-bit VLAN ID.
Resequencing Downstream Service Identifier	A downstream service identifier for which the CMTS signals packet resequencing attributes.

4 ABBREVIATIONS AND ACRONYMS

This specification uses the following abbreviations: ⁷

AC	Attachment Circuit
AGI	Attachment Group Identifier
AS	Autonomous System
ASN	Autonomous System Numbers
ASCII	American Standard Code for Information Interchange
ARP	Address Resolution Protocol
ATM	Asynchronous Transfer Mode
BEB	Backbone Edge Bridge
BGP	Border Gateway Protocol
BPI	Baseline Privacy Interface
BPKM	Baseline Privacy key Management
BSoD	Business Services over DOCSIS
CCM	Continuity Check Message
CE-ID	Customer Edge Identifier
CFI	Canonical Format Indicator
CM	Cable Modem
CMCI	Cable Modem to Customer Premises Equipment Interface
CMIM	CM Interface Mask
CMTS	Cable Modem Termination System
CPE	Customer Premises Equipment
CRC	Cyclic Redundancy Check
DA	Destination Address
DAC	DEMARC Automatic Configuration
DEI	Drop Eligibility Indicator
DEMARC	DPoE Demarcation Device
DHCP	Dynamic Host Configuration Protocol
DIME	Downstream IP Multicast Encryption
DMAC	Destination MAC
DMM	Delay Measurement Message
DMR	Delay Measurement Reply
DOCSIS	Data-Over-Cable Service Interface Specifications
DPoE	DOCSIS Provisioning of EPON
DS	Downstream

⁷ Section modified per L2VPN-N-07.0471-2 on 7/5/07, L2VPN-N-07.0504-1, and L2VPN-N-07.0538-2 on 11/9/07 by KN. Additional revisions per L2VPN-N-12.1065-11 on 9/5/12 by JB. Revised per L2VPN-N-14.1135-5 on 3/27/14 by JB, and L2VPN-N-14.1186-1 on 1/16/15 by PO.

DSCP	Differentiated Services Code Point
DSG	DOCSIS Set-top Gateway
DSID	Downstream Service Identifier
DSx	Dynamic Service addition, change or deletion
DUT	Downstream Unencrypted Traffic
eCM	Embedded Cable modem [eDOCSIS]
eMTA	Embedded Multimedia Terminal Adapter [PKT-PROV]
EPON	Ethernet Passive Optical Network
eRouter	Embedded DOCSIS Router
eSAFE	Embedded Service/Application Functional Entity [eDOCSIS]
eSTB	Embedded Set-top Box
ETH-RDI	Ethernet Remote Defect Indication
EXP	Experimental bits
FD	Frame Delay
FDB	Forwarding Database
FEC	Forwarding Equivalence Class
FLM	Frame Loss Measurement
IANA	Internet Assigned Numbered Authority
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Protocol
IP	Internet Protocol
ISDN	Integrated Services Digital Network
ISP	Internet Service Provider
ITU	International Telecommunication Union
ITU-T	Telecommunication Standardization of the International Telecommunication Union
GMAC	Group Media Access Control
L2	Layer 2
L2CP	Layer 2 Control Protocol
L2PDU	Layer 2 Protocol Data Unit
L2TP	Layer 2 Tunneling Protocol
L2VPN	Layer 2 Virtual Private Network
LAN	Local Area Network
LBM	Loopback Message
LBR	Loopback Reply
LDP	Labeled Distribution Protocol
LLC	Logical Link Control

LSB	Least Significant Bit
LSP	Label Switched Paths
LTM	Linktrace Message
LTR	Linktrace Reply
M	Mandatory
MA	Maintenance Association[802.1Q]
MAC	Media Access Control
MAID	Maintenance Association Identifier [802.1Q]
MD	Maintenance Domain [802.1Q]
MEG	Maintenance Entity Group
MEP	Maintenance association End Point [IEEE 802.1Q]
MEPID	Maintenance Entity Group End Point Identifier
MIB	Management Information Base
MIC	Message Integrity Check
MIP	Maintenance domain Intermediate Point [802.1Q]
MP-BGP	MultiProtocol BGP
MPLS	Multiprotocol Label Switching
MSB	Most Significant Bit
MSO	Multiple Systems Operator
MTA	Multimedia Terminal Adapter
MTU	Maximum Transmission Unit
MULPI	MAC and Upper Layer Protocols Interface
NA	Not Applicable
NLRI	Network Layer Reachability Information
NSI	Network System Interface
OAM	Operations, Administration, and Maintenance
OSPF	Open Shortest Path First
OSSI	Operations System Support Interface
P2P	Point to Point
PB	Provider Backbone
PBB	Provider Backbone Bridging
PCP	Priority Code Point
PDU	Protocol Data Unit
PE	Provider Edge
PICS	Protocol Implementation Conformance Statement
PSN	Packet Switched Network
PW	Pseudowire
PWid	Pseudowire identifier

QoS	Quality of Service
RC	Read-Create
RD	Route Distinguisher
RF	Radio Frequency
RIP	Routing Information Protocol
RO	Read-Only
RSP	Response
RSVP-TE	Resource Reservation Protocol-Traffic Engineering
RT	Route Target
SA	Source Address
SA-MAP	Security Address Mapping
SAID	Security Association Identifier
SAII	Source Attachment Individual Identifier
SF	Service Flow
SID	(Upstream) Service Identifier
SG	Serving Group
SLA	Service Level Agreements
SLM	Synthetic Loss Message
SLR	Synthetic Loss Reply
SMAC	Source MAC
SNMP	Simple Network Management Protocol
SOAM	Service Operations, Administration, and Maintenance
TAII	Target Attachment Individual Identifier
TCI	Tag Control Information
TCP	Transmission Control Protocol
TFTP	Trivial File Transfer Protocol
TLS	Transparent LAN Service
TLV	Type/Length/Value
TOD	Time of Day
TOS	Type of Service
TPID	Tag Protocol Identifier
UCA	User Customer Address
UDC	Upstream Drop Classifier
UNI	User Network Interface
VE	VPLS Edge
VE-ID	VPLS Edge Identifier
VID	VLAN Identifier
VPLS	Virtual Private LAN Services

VPN	Virtual Private Network
VPWS	Virtual Private Wire Service
VSI	Virtual Switch Instance

5 THEORY OF OPERATION (INFORMATIVE)

5.1 L2VPN Features

The ability to implement Layer 2 Virtual Private Networking to arbitrary sets of CMs enables a number of significant DOCSIS features:

- Transparent LAN Service
- Multiple ISP L2VPNs
- Management L2VPNs

5.1.1 Transparent LAN Service⁸

Data networking between the multiple sites of commercial businesses represents a significant business opportunity for cable operators. Commercial data networks are usually implemented with private point-to-point data connections such as Frame Relay, ISDN, or ATM virtual circuits, often with equipment that provides transparent delivery of layer 2 Ethernet LAN packets. A service that interconnects subscriber enterprise LANs with Layer 2 forwarding is called Transparent LAN Service (TLS).

The DOCSIS RFI standard [DOCSIS MAC] was originally intended for residential subscriber connection to the public Internet. This specification standardizes, within DOCSIS, the control and data plane operation of CMTSs and CMs in order to offer Transparent LAN Service to commercial subscriber enterprises.

The term TLS refers to a particular service offering to commercial enterprise customers. The particular technology that provides this service is called a Layer 2 Virtual Private Network (L2VPN). A cable operator offers TLS by implementing one L2VPN for each commercial enterprise.

An example DOCSIS-based commercial TLS service is depicted in Figure 5–1.




Figure 5–1 - Transparent LAN Service

⁸ Section modified per L2VPN-N-07.0471-2 on 7/5/07 by KN and per L2VPN-N-10.0918-2 on 5/10/10 by JB.

Figure 5–1 depicts a Transparent LAN Service offered to two commercial enterprises; one denoted as L2VPN 17, and one as L2VPN 18. All of the CMTSs have the usual set of residential subscribers, which are depicted on CMTS 1 only. CMTS 1 provides L2VPN service to two CMs on L2VPN 17, and one on L2VPN 18. CMTS 2 provides L2VPN service to one L2VPN 18 CM. CMTS 3 provides service to one CM on L2VPN 17, and one on L2VPN 18.

The example shows that the cable operator's L2 backbone implements a single Virtual LAN (VLAN) for each customer. In this document, the term VLAN has a specific meaning as referring to the [802.1Q] definition, as a subset of LANs, within a Bridged Network, to which is assigned a 12-bit VLAN ID. In this example, CMTS 1 directly encapsulates upstream L2 packets from L2VPN 17 onto an IEEE 802.1Q tagged Ethernet packet with a VLAN ID tag 17, and forwards them onto a trunk Ethernet Network System Interface (NSI) port, to the cable operator's backbone.

In the example, CMTS 1 implements multipoint L2 Forwarding, so it is responsible for bridging packets between its two CMs attached to L2VPN 17. Bridging requires the CMTS to learn the Source MAC (SMAC) addresses of CPE17A and CPE17B, and associate them with the CM to which each CPE device is attached.⁹

CMTS 1 implements only a single attachment circuit to VLAN 17 on the backbone. When a downstream unicast packet from VLAN 17 arrives at CMTS 1, it looks up the Destination MAC (DMAC) in its learning database and forwards the packet to the correct CM.

CMTS 3, however, may implement only point-to-point L2 forwarding, where it transparently forwards all individual and group MAC destined packets, in a point-to-point manner, between the CM attached to CPE 17C and IEEE 802.1Q VLAN ID 17 on its NSI Ethernet port, to the backbone.

In the backbone, a cable operator Layer 2 Bridge connects the various Ethernet trunk interfaces from the CMTSs, and bridges together each VLAN. The TLS service offered by the operator to L2VPN17 provides for a transparent layer 2 bridged connection between CPEs 17A, 17B, and 17C. From the enterprise customer's point of view, such CPE are managed and operated as if they were on a customer-private Ethernet LAN. Usually, they will have an IP address on the same IP subnet owned by the enterprise. The enterprise usually assigns the IP address to each CPE, and typically has its own DHCP server to do so. Indeed, each enterprise can use the same overlapping private IP subnet space. Unlike Layer 3 VPN technologies, the cable operator does not need to co-ordinate IP address subnet assignment with the enterprise customers. From the operator's point of view, the enterprise LAN subscribers are completely isolated at layer 2 from all other residential subscribers, and from every other L2VPN.

An enterprise TLS may include not only the LANs attached to CMs, but also any other LANs bridged to the customer's VLAN in the IEEE 802.1Q-compliant Bridge in the cable operator's backbone.

As shown in Figure 5–2, cable operators can offer both port-based Ethernet Private Line (EPL) and VLAN-based Ethernet Virtual Private Line (EVPL) service over an MPLS network. In such a network, the CMTS can serve as a Provider Edge (PE) router and map DOCSIS service flows into MPLS pseudowires.

In the example below, a cable operator provides EPL service to a single commercial enterprise denoted as L2VPN 17. Both CMTSs have the usual set of residential subscribers, which are depicted on CMTS 1 only. CMTS 1 provides L2VPN service to two CMs on L2VPN 19. CMTS 2 provides L2VPN service to one L2VPN 19 CM.

The example shows that the cable operator's MPLS backbone implements a single pseudowire for each customer. In this example, CMTS 1 directly encapsulates upstream L2 packets from L2VPN 19 into an IETF [RFC 4448] Ethernet over MPLS pseudowire and forwards them to the cable operator's backbone via a Network System Interface (NSI) port.

In the example, CMTSs 1 and 2 implement point-to-point L2 Forwarding. Each CMTS implements only a single attachment circuit to the IP/MPLS pseudowire on the backbone. When a downstream unicast packet from VLAN 19 arrives at CMTS 1, it looks up the pseudowire mapping in its database and transparently forwards all individual and group MAC destined packets, in a point-to-point manner, to the correct CM.

In the backbone, cable operator core IP/MPLS routers forward pseudowire traffic between the two MPLS PE CMTSs. The TLS service offered by the operator to L2VPN 19 provides for a transparent layer 2 bridged connection between CPEs 19A and 19B. From the enterprise customer's point of view, such CPE are managed and operated as if they were on a customer-private Ethernet LAN. As with the example in Figure 5–1, from the operator's point of

⁹ Paragraph modified per L2VPN-N-07.0592-1 on 1/24/08 by KN.

view, the enterprise LAN subscribers are completely isolated at layer 2 from all other residential subscribers, and from every other L2VPN.

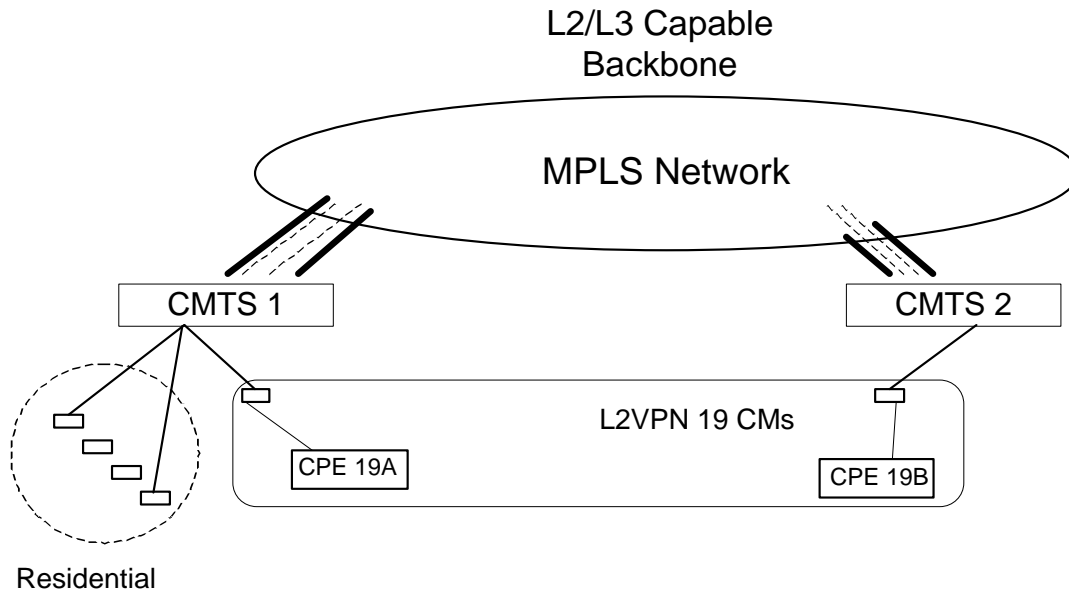


Figure 5–2 - TLS over a MPLS Network

5.1.2 Multiple ISP L2VPNs

The L2VPN feature permits a cable operator to support multiple Internet Service Providers (ISPs) by providing a separate L2VPN for each ISP. The cable operator provides all CM provisioning, and the CM configuration file determines a L2VPN for forwarding all CPE traffic. Each ISP is assigned a separate L2VPN. The ISP is responsible for providing the DHCP servers and IP addressing for all CPEs on the CMs attached to their L2VPN.

The advantage of L2VPNs for multiple ISP operation is that it completely separates the IP address space management and IP routing of the ISP, from that of the cable operator. In contrast, multiple ISP features, based on layer 3 VPNs, usually require co-ordination of IP address assignment and router security configuration between the MSO Provider Edge and ISP Customer Edge routers.

5.1.3 Management L2VPNs

The DOCSIS L2VPN feature allows a CMTS to implement an L2VPN solely for the provisioning and management of embedded Cable Modems (eCMs) and embedded Service/Application Functional Entities (eSAFEs) [eDOCSIS], such as an embedded Multimedia Terminal Adapter (eMTA) [PKT-PROV], or eRouter [eRouter].¹⁰ Implementing a separate L2VPN for provisioning and management of eCM and eSAFE traffic isolates those devices from the Internet and from the subscriber, enhancing security.

Prior to registration, the CM transmits on a temporary SID, and all such traffic is considered to be forwarded by the Non-L2VPN Forwarder. A CMTS could be implemented to forward pre-registration traffic onto a single Provisioning L2VPN. The Provisioning VPN would be configured in a vendor-specific manner.

When a CM registers, it reads L2VPN Encodings from its configuration file that can configure its eCM and eSAFE devices to forward-on an L2VPN. This post-registration L2VPN is called a Management L2VPN, because post-registration traffic is primarily SNMP for managing the device.

¹⁰ Revised per L2VPN-N-12-1065-11 on 9/5/12 by JB.

5.1.4 Other L2VPN-enabled Features

Some features required by the specification are enhancements to overall DOCSIS operation otherwise unrelated to Layer 2 VPNs:

- Interface-based Classification
- DUT Filtering
- Enabling eSAFE DHCP Snooping Control

Interface-based classification allows packets to be classified according to the CM internal or external bridge port interface, and is described in Section 6.6.5. This feature may be used, for example, to classify packets to, or from the embedded MTA interface, without relying on the particular IP subnet of that interface.

Downstream Unencrypted Traffic (DUT) filtering is applicable to CMTS vendor-specific Layer 3 VPN operation to prevent the Group MAC traffic that is broadcast to residential CMs from leaking into the supposedly private CPE networks of Layer 3 VPN subscribers. DUT Filtering is described in Section 6.5.2.1.

DHCP Snooping Control is an explicit TLV that authorizes the CMTS to automatically learn the MAC address of embedded eSAFE hosts, such as eMTAs, by intruding on their DHCP traffic. This may be used in conjunction with CMTS vendor-specific features that forward DHCP or other packets from eSAFE hosts in a special manner. The Enable eSAFE DHCP Snooping Control feature is described in Section 6.6.4.1.

5.1.5 High Availability ¹¹

This document uses the term "Layer 2 trunk port" to describe an NSI port configured to forward L2VPN packets encapsulated with a layer 2 tag as specified by IEEE 802, e.g., IEEE 802.1Q.

This specification defines requirements for the CMTS to maintain L2VPN forwarding in the event of any single NSI port failure. One technique is to support forwarding of all L2VPN traffic using a layer 2 IEEE 802 Ethernet encapsulation (e.g., 802.1Q) through a single "layer 2 trunk port" at a time, and support configuration of multiple possible trunk ports, only one of which is "active". Another technique is to implement L2VPN forwarding across a set of NSI ports using [802.1AX] Link Aggregation.

¹¹ Added section per L2VPN-N-12.1065-11 on 9/5/12 by JB, L2VPN-N-14.1186-1 on 1/15/15 by PO.

5.2 CMTS Layer 2 Forwarding Architecture

5.2.1 L2VPN and Non-L2VPN Forwarding

A CMTS is considered to have an entirely separate packet forwarder for L2VPN forwarding that differs from Non-L2VPN forwarder for residential traffic, as depicted in Figure 5–3.

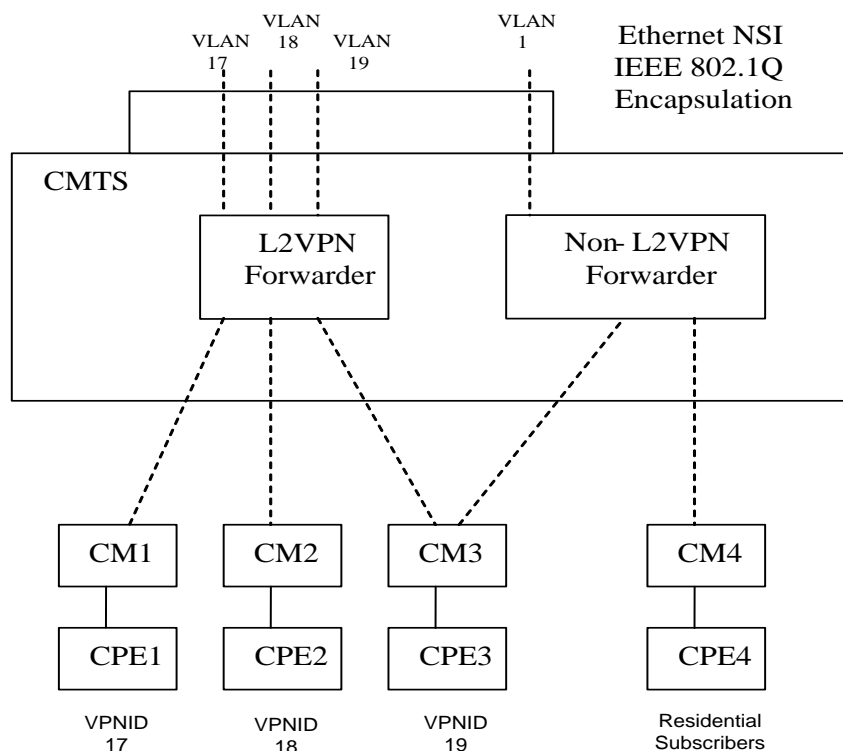


Figure 5–3 - CMTS L2VPN and Non-L2VPN Forwarding

To support L2VPN operation, a CMTS's Network System Interface (NSI) has to be capable of distinguishing L2VPN from non-L2VPN downstream traffic, and determining the L2VPN of the downstream traffic.¹² The encapsulation format of L2VPN traffic on a CMTS' NSI ports and the particular field values, within that encapsulation that distinguish a particular L2VPN, are called the NSI L2VPN Encapsulation information. In the example above, IEEE 802.1Q VLAN ID tags are used as the L2VPN Encapsulation format on an Ethernet NSI port.

In general, L2VPN and non-L2VPN traffic is mixed on the same NSI port. In the example above, the CMTS implements a non-L2VPN IP router interface on VLAN ID 1, which may even be the native VLAN, with an untagged encapsulation. Residential traffic, such as CPE4 connected to CM4, continues to be routed through the CMTS's IP routing forwarder onto the router's sub-interface on VLAN ID 1. The other CPE, however, are bridged at layer 2 from the Ethernet interface of the CM to a configured 802.1Q VLAN ID on the NSI port. In Figure 5–3 above, the CMTS implements a Point-to-Point forwarding model where it forwards CPE traffic from CM1 to 802.1Q VLAN ID 17, CPE traffic from CM2 to 802.1Q VLAN ID 18, and CPE traffic from one of the upstream service flows of CM3 to 802.1Q VLAN ID 19. The other upstream service flow of CM3 is forwarded to the non-L2VPN forwarder.

¹² Revised sentence per L2VPN-N-12.1065-11 on 9/5/12 by JB.

In the upstream direction, the CMTS distinguishes L2VPN from non-L2VPN traffic based on the Upstream Service Flow from which the traffic arrives, and the source MAC address of the traffic. Certain Upstream SFs are configured with Forwarding L2VPN Encodings that identify a particular L2VPN. The L2VPN encoding includes a CM Interface Mask (CMIM) that identifies which CM-side hosts forward upstream to the L2VPN. By default, only CPE hosts attached to the CMCI interface of a CM forward to an L2VPN; the CM and its internal eSAFE hosts do not forward to an L2VPN.

An SF configured to forward CPE traffic to an L2VPN is considered to be an attachment circuit in the context of IETF Virtual Private LAN Service (VPLS). A CMTS VPLS L2VPN Forwarder is responsible for forwarding packets between attachment circuits and pseudowires on NSI ports (e.g., MPLS or L2TPv3 tunnels).

5.2.2 Point-to-Point and Multipoint L2VPN Forwarding Modes ¹³

This specification uses the term layer 2 forwarding rather than bridging because commercial L2VPN service can be offered without necessarily implementing a learning MAC Layer Bridge on the CMTS as defined by [802.1Q]. The CMTS MAY implement a Point-to-Point layer 2 forwarding mode that forwards packets between a single NSI port and a single CM (or SF). If the CMTS does implement a learning MAC layer bridge between NSI and RF interfaces, this specification terms it the Multipoint layer 2 forwarding mode.

In Point-to-Point L2VPN Forwarding Mode, each attachment circuit has a different NSI Encapsulation value. For example, with IEEE 802.1Q encapsulation, each attachment circuit (i.e., CM or SF) is configured with a different 802.1Q VLAN ID. Similarly, using point-to-point Ethernet over MPLS [RFC 4385] [RFC 4448] encapsulation, each attachment circuit is configured with a different pseudowire identifier. In Point-to-Point mode, the L2VPN forwarder simply forwards upstream and downstream data between one NSI port and one attachment circuit, without learning the MAC addresses of CPE packets. The logical VPNID to which a CM or SF attaches should be configured with the attachment circuit, but its value is otherwise ignored by the CMTS in Point-to-Point forwarding mode. An external L2VPN Bridge on the cable operator's backbone actually performs the layer 2 MAC address learning for each L2VPN, and bridges packets between the VLAN IDs or pseudo-wires of the packets within their NSI Encapsulation.

¹³ Revised per L2VPN-N-10.0918-2 on 5/10/10 by JB.

An example of Point-to-Point Forwarding Mode is depicted in Figure 5–4.

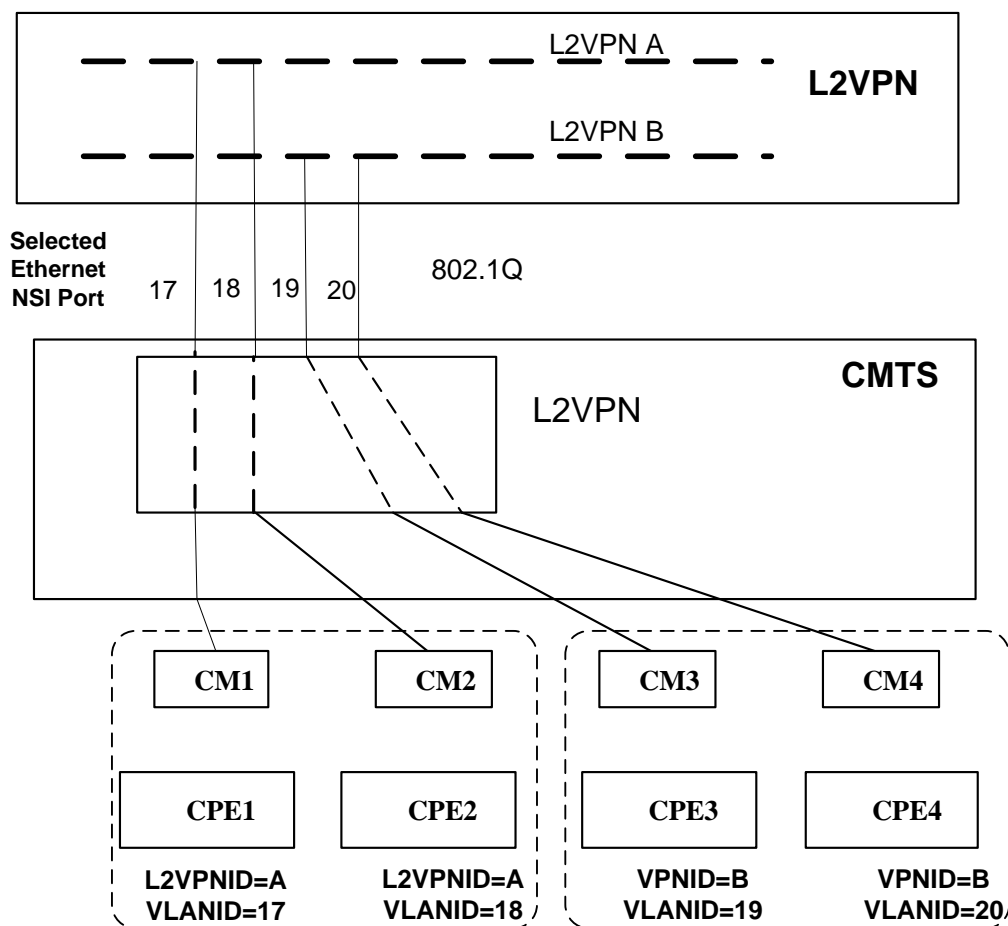


Figure 5–4 - Point-to-Point Forwarding Mode

Four CMs are configured for L2VPN operation. Each CM's L2VPN Encoding includes a logical VPNID A or B, along with a statically configured NSI Encapsulation subtype that configures the use of [802.1Q] with a different VLAN ID for each CM (VLAN IDs 17 through 20). The CMTS's L2VPN Forwarder forwards traffic from the NSI port on those VLAN IDs in a point-to-point manner, to and from the configured CM. Although the L2VPN Forwarder in Point-to-Point mode does not use the VPNID configuration, it still has to be configured in each forwarding L2VPN Encoding for at least information purposes.¹⁴ An L2VPN Bridge Layer 2 switch, external to the CMTS, is configured to treat the NSI encapsulations for VLAN ID 17 and 18 as separate logical bridge ports for L2VPN A, and to learn CPE MAC addresses on those bridge ports. Likewise, the external L2VPN Bridge is configured to consider the encapsulations with VLAN IDs 19 and 20, as separate bridge ports of the broadcast domain that is L2VPN B.

With IEEE 802.1Q NSI Encapsulation Point-to-Point Forwarding Mode, the number of L2VPN subscriber modems supported on a CMTS is limited to 4093 CMs, due to the 12-bit limit of an IEEE 802.1Q VLAN ID.¹⁵

¹⁴ Revised sentence per L2VPN-N-12.1065-11 on 9/5/12 by JB.

¹⁵ Paragraph revised per L2VPN-N-07.0461-2 on 7/5/07 by KN.

Multipoint Forwarding mode means that the CMTS forwards downstream L2VPN packets to potentially multiple cable modems. The CMTS builds a layer 2 Forwarding Database (FDB) of the CPE MAC addresses it learns from the source MAC address of upstream packets. A Multipoint L2VPN Forwarder uses this FDB to select which CM to forward downstream L2VPN traffic. If the destination is a group MAC address, or is an unknown individual MAC address, a Multipoint L2VPN Forwarder floods the traffic to all attachment circuits and NSI ports other than the one from which the packet was received. A Multipoint L2VPN forwarder also directly forwards packets between attachment circuits (CMs or SFs), configured to the same logical L2VPN.

With Multipoint Forwarding, an NSI Encapsulation value is needed only for each logical L2VPN, not for each attachment circuit. This allows support of any number of *modems* for L2VPN service, because the 12-bit IEEE 802.1Q VLAN ID, used as an NSI Encapsulation value, will only limit the number of enterprise L2VPN *networks*.

An example of Multipoint Forwarding Mode is depicted in Figure 5–5.




Figure 5–5 - Multipoint L2VPN Forwarding Mode Example

In this example, both CM1 and CM2 are configured for L2VPN forwarding to VPNID A, and configured to use an NSI Encapsulation of IEEE 802.1Q VLAN ID 17. The Multipoint L2VPN forwarder learns the MAC addresses of CPE1 and CPE2 in order to determine to which CM to forward downstream unicast traffic received from the network port on VLAN ID 17. Likewise CM3 and CM4 are configured with VPNID B and both are configured to use IEEE 802.1Q VLAN ID 18 as their NSI Encapsulation. The Multipoint L2VPN forwarder learns the MAC addresses of CPE3 and CPE4 on L2VPN B. The non-L2VPN forwarder is not shown in the Figure 5–5.

This specification permits qualification of CMTSs with either Point-to-Point or Multipoint forwarding modes. DOCSIS qualification testing uses the forwarding mode indicated by the vendor's PICS submission for all L2VPNs.¹⁶ This specification is written assuming that a CMTS selects one mode or the other for all L2VPNs. There are no requirements, however, that prevent a vendor from implementing different forwarding modes for different sets of L2VPNs.

¹⁶ Revised sentence per L2VPN-N-12.1065-11 on 9/5/12 by JB.

6 L2VPN OPERATION (NORMATIVE)

6.1 CMTS Bridging Model Requirements¹⁷

The CMTS MUST transparently forward DOCSIS L2PDUs received from an Upstream Service Flow configured to receive packets for a particular L2VPN to NSI ports configured to encapsulate packets for that L2VPN. The CMTS MUST transparently forward packets received with an NSI encapsulation configured for a particular L2VPN to a downstream DOCSIS L2PDU encrypted in a SAID unique to that L2VPN and CM to which the packet is forwarded.

A CMTS SHOULD implement a VLAN-capable bridging function as specified by [802.1Q]. For purposes of [eRouter] conformance, each bridge port implemented on an RF interface SHOULD be considered to be a fully-tagged 802.1Q interface for which the incoming VLAN ID is determined by the upstream SID, and the outgoing VLAN ID is tagged with a BPI SAID. A L2VPN compliant CMTS MUST NOT insert an 802.Q tag on downstream RF packets.¹⁸

The CMTS MAY restrict configuration of an NSI Encapsulation service multiplexing value (e.g., IEEE 802.1Q VLAN ID) to a single SF. In this Point-to-Point forwarding mode, the CMTS MAY omit learning of CPE MAC addresses into a Forwarding Database. A Point-to-Point CMTS MUST support multiple per-SF L2VPN Encodings with the same NSI Encapsulation subtype as long as they are on the same CM. SFs carrying L2VPN traffic could be carried over bonded channels if the CMTS and CM support Channel Bonding as defined in [MULPIv3.0]. Consequently, the use of resequencing DSIDs is not precluded by this specification.

If the CMTS permits more than one SF to be configured to bridge to the same NSI Encapsulation service multiplexing value, it is said to implement Multipoint forwarding mode. In Multipoint forwarding mode, the CMTS MUST associate learned CPE source MAC addresses with the particular CM from which they were learned.

A CMTS MUST support both L2VPN and non-L2VPN forwarding on the same RF MAC domain. A CMTS MUST transparently bridge CPE traffic from CMs configured with L2VPN Encodings according to this specification. A CMTS MUST forward with its normal, non-L2VPN packet forwarding algorithms CPE traffic from CMs with no L2VPN Encodings, except as specified in this specification.

A CMTS MUST support both L2VPN and non-L2VPN forwarding of upstream traffic from different service flows when only per-SF L2VPN Encodings are signaled. This traffic could arrive on a single channel or on a bonded channel group, provided that Channel Bonding is supported by both the CM and CMTS.

6.2 Configuring L2VPN Forwarding¹⁹

A set of one or more L2VPN Encoding configuration settings in a CM configuration file controls whether and how the CMTS performs L2VPN forwarding of upstream and downstream CPE packets.

The L2VPN Encoding parameter is encoded as a General Extension Information (GEI) parameter, meaning it is encoded as a subtype of the Vendor Specific Information type 43 parameter using, vendor ID 0xFFFFFFFF (DOCSIS MAC). By encoding the L2VPN Encoding as a GEI parameter, it may be included in the configuration file of any DOCSIS CM, including DOCSIS 1.0 CMs.

The L2VPN Encoding parameter may appear in the following locations:

- At the top level of a CM configuration file, in which case it is called a per-CM L2VPN Encoding;
- As a subtype of a GEI nested in an Upstream Service Flow Encoding (type 24), in which case it is called a per-SF or Forwarding L2VPN Encoding;
- As a sub-type of a GEI nested in a Downstream Packet Classification Configuration Setting (type 23), in which case it is called a Downstream Classifier L2VPN Encoding;

¹⁷ Section modified per L2VPN-N-07.0504-1 on 11/9/07 by KN.

¹⁸ Section modified per L2VPN-N-07.0461-2 on 7/5/07 by KN.

¹⁹ Section modified per L2VPN-N-07.0471-2 on 7/5/07 by KN and revised per L2VPN-N-10.0918-2 on 5/10/10 by JB.

- As a sub-type of a GEI nested in an Upstream Packet Classifier Configuration Setting (type 22), in which case it is called an Upstream Classifier L2VPN Encoding.

The L2VPN Encoding parameter itself is defined as a multi-part parameter with several nested subtype parameters. The following Table 6–1 lists each of the subtypes and describes in which location the subtype is defined, as required or optional for that location.

Table 6–1 - L2VPN Encoding Subtype Location Summary²⁰

Subtype Number	Subtype Parameter	Top Level (per-CM)	Upstream Service Flow	Downstream Classifier	Upstream Classifier
43.5.1	VPN Identifier	Required	Required	Required	
43.5.2	NSI Encapsulation	Optional(3)			
43.5.3	Enable eSAFE DHCP Snooping	Optional(1)			
43.5.4	CM Interface Mask	Optional		Optional(1)	Optional(1)
43.5.5	Attachment Group ID	Optional(3)			
43.5.6	Source Attachment Individual ID	Optional(3)			
43.5.7	Target Attachment Individual ID	Optional(3)			
43.5.8	Upstream User Priority		Optional		
43.5.9	Downstream User Priority Range			Optional	
43.5.10	L2VPN SA-Descriptor	Required(2)			
43.5.11	deprecated				
43.5.12	deprecated				
43.5.13	L2VPN Mode (4)				
43.5.14	TPID Translation (4)				
43.5.15	L2CP Processing (4)				
43.5.16	Demarc Autoconfiguration (4)				
43.5.17	deprecated (4)				
43.5.18	Pseudowire Class (4)				
43.5.19	Service Delimiter (4)				
43.5.20	Virtual Switch Instance (4)	Optional			
43.5.21	BGP Attribute	Optional	Optional	Optional	
43.5.22	L2VPN Serving Group (4)				
43.5.23	Pseudowire Signaling	Optional			
43.5.24	Service OAM	Optional			
43.5.25	Network Timing Profile Reference (4)				
43.5.26	L2VPN DSID (5)	Optional			
43.5.43	Vendor-Specific	Optional	Optional	Optional	Optional

Table Notes:

- (1) The CMTS MUST accept a parameter identified as optional (1) in Table 6–1 in a non-forwarding L2VPN Encoding.
- (2) The CMTS inserts the L2VPN SA-Descriptor Subtype in its first message to a CM in any MAC Management Message that includes a Forwarding L2VPN Encoding; the L2VPN SA-Descriptor Subtype is not configured in a CM configuration file.
- (3) This is a Per L2VPN configuration on the NSI port that is defined in a per-CM L2VPN Encoding only for Point-to-Point forwarding mode.
- (4) This TLV is specific to DOCSIS Provisioning of EPON (DPoE) Systems.
- (5) This TLV is specific to DOCSIS 3.0 and later L2VPN systems

²⁰ Revised table per L2VPN-N-12.1065-11 on 9/5/12 by JB.

If a subtype is not defined as Required or Optional in a location, the CMTS SHOULD silently ignore it when it appears in that location. If a subtype is not defined as Required or Optional in a location, the cable modem SHOULD silently ignore it when it appears in that location. A CMTS MUST silently ignore unrecognized subtypes in an L2VPN Encoding. A CM MUST silently ignore unrecognized subtypes in an L2VPN Encoding.

The top-level L2VPN Encoding controls the per-L2VPN CM and CMTS behavior specific to a particular L2VPN. The Upstream Service Flow L2VPN Encoding specifies which upstream service flow(s) will carry L2VPN traffic. Proper L2VPN operation requires at least one upstream service flow to be configured for L2VPN forwarding.

Because multiple upstream service flows can be configured to forward to the same L2VPN, all of the per-L2VPN parameters common to the L2VPN itself are encoded in the single top-level L2VPN encoding rather than requiring or permitting them to be duplicated in multiple upstream service flow encodings.

Upstream L2VPN forwarding is configured on a per-SF basis. The cable operator can configure at least one upstream service flow in a CM configuration file with an L2VPN Encoding that defines the VPN Identifier to which the CMTS forwards upstream traffic from that SF. The per-CM or top-level L2VPN encoding is required in a CM configuration file only for point-to-point forwarding mode, in order to define the NSI encapsulation format for that L2VPN so that the CMTS can determine to which CM to forward downstream L2VPN traffic. In point-to-point forwarding mode, a CMTS MUST support one per-CM L2VPN Encoding with an NSI Encapsulation subtype and at least one Upstream Service Flow L2VPN Encoding per CM. With Multipoint forwarding mode for an L2VPN, multiple CMs may forward to multiple NSI encapsulations, so any per-CM NSI encapsulation configuration is not defined.

The simplest CM configuration file for L2VPN operation contains:

- For Multipoint mode, a single per-SF L2VPN Encoding within the primary upstream SF definition; or
- For Point-to-Point mode, single per-SF L2VPN Encoding within the primary upstream SF definition and a single per-CM L2VPN Encoding with an NSI Encapsulation subtype for that L2VPN.

In a Registration Response message, the CMTS always includes a per-CM L2VPN encoding (adding a per-CM L2VPN encoding if necessary) that provides at least one L2VPN SA-Descriptor for encrypting and labeling downstream packets as L2VPN traffic for the CM. The CMTS MAY assign more than one SAID to the same L2VPN, in which case multiple L2VPN SA-Descriptor subtypes may appear in a top-level L2VPN Encoding.

Unless configured otherwise, the CMTS delivers downstream L2VPN traffic to a single cable modem on the CM's primary downstream service flow. The operator can specify enhanced Quality of Service (QoS) for downstream L2VPN traffic with a separate downstream service flow for L2VPN forwarding in the CM's configuration file. Downstream L2VPN traffic can be classified to that particular downstream service flow by defining a classifier that includes a Downstream Classifier L2VPN Encoding that references the service flow.

The CMTS MUST reject the registration of a CM with an invalid L2VPN Encoding. A valid CM configuration contains any number of per-SF L2VPN Encodings, Downstream Classifier L2VPN Encodings and Upstream Classifier L2VPN Encodings. The CMTS MUST accept a CM registration request that contains multiple per-SF L2VPN Encodings that forward to the same VPN ID.

A valid per-SF L2VPN Encoding appears as a subtype in the Upstream Service Flow Encoding (type 24) of a DOCSIS 1.1 CM configuration file, REG-REQ, REG-REQ-MP, DSA-REQ, or DSC-REQ message. The per-SF L2VPN Forwarding Encoding configures the CMTS to perform L2VPN bridge forwarding for all CPE packets received in the described service flow. A valid per-SF L2VPN Forwarding Encoding contains one L2VPN ID subtype. The CMTS includes a per-CM L2VPN encoding in its REG-RSP or REG-RSP-MP. After registration, a CM MAY include per-CM L2VPN encodings at the top level of Dynamic Service MAC Managements messages that otherwise add, change, or delete forwarding per-SF L2VPN encodings.

In order to configure particular CPE MAC addresses for L2VPN forwarding, the CM can be configured with Upstream Packet Classification Encodings that match the desired source CPE MAC address. The CM classifies the packet to an upstream service flow that is configured to forward to a particular L2VPN. The Upstream Packet Classification Encoding that references a forwarding Upstream SF L2VPN encoding does not contain a VPNID subtype itself.

The CMTS MUST consider an upstream service flow to be configured for per-SF L2VPN forwarding when a REG-REQ, REG-REQ-MP, DSA-REQ, or DSC-REQ contains exactly one valid per-SF L2VPN Forwarding Encoding within an Upstream Service Flow Encoding. A valid per-SF L2VPN Forwarding Encoding contains one VPNID subtype. The CMTS MUST reject a service flow transaction that contains more than one per-SF L2VPN Encoding.

The CMTS MUST accept a valid DSC-REQ with a valid per-SF L2VPN Encoding and change the upstream forwarding treatment of packets received from that SF accordingly. This includes, for example, adding, changing, or deleting any permitted subtype of a per-SF L2VPN encoding, including the VPN ID subtype.

The CMTS MUST remove per-SF L2VPN forwarding for an SF when the SF is deleted with a valid Dynamic Service Delete (DSD) transaction or a Dynamic Service Change (DSC) transaction completes that omits a previously signaled per-SF L2VPN Encoding.

The CMTS MUST support multiple per-SF L2VPN Encodings, each on a separate SF, with the same VPNID Subtype value.

A multipoint forwarding CMTS MAY accept the per-VPN subtypes defined only for point-to-point mode, but CMTS operation with different subtype values on different CMs is not defined.

6.2.1 VPNID Subtype²¹

The VPNID Subtype is an opaque byte sequence that identifies a logical Layer 2 Virtual Private Network. All hosts attached to the same logical L2VPN communicate with each other as if they were attached to the same private LAN. An L2VPN is a network that forwards packets based only on layer 2 information such as the Ethernet MAC addresses and any VLAN ID tags encapsulating the packet. The term VLAN ID should be used only to describe the 12-bit VLAN ID field encoded in an IEEE 802.1Q tag of an L2VPN packet forwarded on an NSI port.

A cable operator is expected to configure a unique VPNID for each commercial enterprise to which it offers Transparent LAN service. The cable operator may choose any format desired for the VPNID, but it should be globally unique. One suggested approach is [RFC 2685], which defines a mechanism for assigning 7-byte globally unique VPN IDs, based on combining a 3-byte Organization Unique ID for the organization assigning the ID (e.g., the cable operator itself), with a 4-byte VPN ID assigned by that organization. Another suggested approach is [RFC 4364], which describes an 8-byte Route Distinguisher that may be used as a globally-unique VPNID.

The CMTS MUST ignore a per-SF L2VPN encoding that omits a VPNID subtype or that contains more than one VPNID subtype. The CMTS MUST support at least four (4) different values of VPNID per CM, signaled in four or more per-SF L2VPN Encodings.

6.2.1.1 BGP VPNID Subtype²²

The BGP VPNID Subtype is a 4-byte integer value that also identifies a logical Layer 2 Virtual Private Network. BGP VPNID Subtype is required only if BGP auto-discovery and LDP signaling [RFC 6074] is to be used by CMTS for L2VPN dynamic discovery.

A cable operator is expected to configure a unique BGP VPNID for each commercial enterprise to which it offers L2VPN service, if the operator intends to use BGP auto-discovery and LDP signaling of pseudowires at the CMTS. A cable operator may configure the same integer value in both VPNID and BGP VPNID Subtypes for operational simplicity. Please see Section 6.2.7.2.2 (BGP Auto-discovery) for more details.

6.2.2 Downstream Classifier L2VPN Encoding²³

A Downstream Classifier L2VPN Encoding is an L2VPN Encoding that appears in a Downstream Packet Classification Encoding [DOCSIS MAC]. The presence of an L2VPN Encoding within a Downstream Packet Classification Encoding restricts the classifier to apply to only packets forwarded by the L2VPN Forwarder. Furthermore, only classifiers that contain an L2VPN Encoding apply to packets forwarded by the L2VPN Forwarder. In other words, downstream classifiers apply either to L2VPN or non-L2VPN traffic, but never both.²⁴

²¹ Revised per L2VPN-N-10.0918-2 on 5/10/10 by JB, L2VPN-N-14.1186-1 on 1/15/15 by PO.

²² Section added per L2VPN-N-12.1065-11 on 9/5/12 by JB.

²³ Section revised per L2VPN-N-12.1065-11 on 9/5/12 by JB.

²⁴ Paragraph modified per L2VPN-N-07.0471-2 on 7/5/07 by KN.

A Downstream Classifier L2VPN Encoding may contain zero or one VPN ID subtypes and/or zero or one Downstream User Priority Range subtypes. It may contain no subtypes at all (i.e., a zero-length 43.5 parameter), in which case the classifier applies to all L2VPN forwarded downstream packets to the CM, regardless of VPN ID or user priority.

The presence of a VPNID subtype within a Downstream Classifier L2VPN Encoding instructs the CMTS to apply the classifier to only downstream L2VPN-forwarded traffic on the indicated L2VPN. Because the L2VPN of L2VPN-forwarded traffic is always *implied* on the DOCSIS RF interface, and is not explicitly present in the packet contents, this is the only way to classify downstream L2VPN-forwarded traffic to a particular service flow based on the VPNID itself.

If the Downstream Classifier L2VPN Encoding contains a Downstream User Priority Range subtype, the classifier applies only to encapsulated L2VPN packets received on the NSI with a user priority within the indicated range (inclusive). This allows high priority L2VPN traffic to be classified to downstream service flows with enhanced QoS.

The user priority value matched by a Downstream User Priority Range subtype is the user priority encoded in the (802.1Q or MPLS or L2TPv3) encapsulation header of the L2VPN packet received by the L2VPN forwarder on the NSI. Please refer to Section 6.7.2.1.1 for user priority description.

A CMTS MUST reject the registration of a CM with a Downstream Packet Classifier Encoding that contains more than one L2VPN Encoding.

6.2.3 L2VPN SA-Descriptor Subtype

The L2VPN SA-Descriptor Subtype is a multiple-part encoding defined in the BPI+ Specification [BPI+] that provides:

- the Baseline Privacy (BPI) Security Association Identifier (SAID) that the CMTS uses to encrypt downstream L2VPN traffic for the L2VPN identified in the L2VPN Encoding;
- a Cryptographic Suite that identifies the encryption algorithm;
- and a Security Association Type (SA-Type).

A CMTS MUST encode the L2VPN SA-Descriptor as a Dynamic (2) SA-Type. A CM MUST ignore the SA-Type and consider it to be of type Dynamic (2).

Upstream L2VPN traffic is always encrypted in the Primary SAID of the CM that transmits the traffic upstream.

The L2VPN SA-Descriptor Subtype is not signaled in a CM configuration file. Instead, the CMTS adds one or more L2VPN SA-Descriptor Subtypes to top-level per-CM L2VPN Encodings of its REG-RSP or REG-RSP-MP to the CM, adding the per-CM L2VPN Encoding to the REG-RSP or REG-RSP-MP if necessary. After a CM completes BPI Authentication, it initiates a Traffic Encrypting Key (TEK) transaction with the CMTS for each L2VPN SA-Descriptor in a REG-RSP or REG-RSP-MP message.²⁵

The CMTS includes an L2VPN SAID in an L2VPN SA-Descriptor Subtype encoding at the top level of a CMTS-initiated DSA-REQ or DSC-REQ message that otherwise defines a L2VPN forwarding upstream SF. Likewise, the CMTS includes an L2VPN SA-Descriptor Subtype in a top-level L2VPN Encoding in its DSA-RSP or DSC-RSP responses to a CM-initiated dynamic service transaction, defining an L2VPN forwarding upstream service flow. A SAID signaled in an L2VPN SA-Descriptor Subtype Encoding is referred to as an L2VPN SAID. A SAID known to the CM only from messages other than an L2VPN SA-Descriptor Subtype is called a non-L2VPN SAID. Section 6.5 below describes how BPI encryption isolates L2VPN and non-L2VPN traffic on the RF network.

After the completion of any Dynamic Service MAC Management message transaction that introduces a new SAID to the CM, the CM initiates a TEK transaction with the CMTS to obtain the keying material for the new SAID.

In Point-to-Point L2VPN forwarding mode, the CMTS assigns an Individual L2VPN SAID to each CM. If the CM forwards more than one L2VPN, the CMTS assigns a different Individual L2VPN SAID for each L2VPN. In

²⁵ Paragraph modified per L2VPN-N-07.0741-2 on 7/5/07 by KN.

Multipoint L2VPN forwarding mode, the CMTS assigns a Group L2VPN SAID for all CMs forwarding the L2VPN to share.

6.2.4 Vendor-Specific L2VPN Encoding

The Vendor Specific L2VPN Encoding subtype is accepted in any L2VPN Encoding location, and provides information specific to the CMTS or CM vendor. For example, it may indicate to a CMTS vendor a particular NSI port sub-interface to which the L2VPN forwards the traffic in a point-to-point model. The Vendor Specific L2VPN Encoding may be binary or ASCII; its definition is left to the CMTS vendor.

A CMTS implementation MAY permit a Vendor Specific L2VPN Encoding to *replace* an otherwise required VPNID or NSI Encapsulation subtype, but vendor-specific L2VPN Encodings MUST NOT be required by a CMTS for L2VPN certification testing.

6.2.5 Configuration Error Requirements

A Multipoint forwarding CMTS MUST reject—with a reject-multipoint-NSI confirmation code—a registration or dynamic service transaction that attempts to configure multiple upstream forwarding L2VPN Encodings to the same L2VPN ID but with different values of the NSI Encapsulation, AGI, TAIL, or SAIL subtypes.

A Point-to-Point forwarding CMTS MUST reject—with a reject-VLAN-ID-in-use confirmation code—a registration or service flow transaction with an L2VPN NSI Encapsulation subtype that requires forwarding on the L2VPN Selected Port with a VLAN ID already assigned for non-L2VPN purposes. A Point-to-Point forwarding CMTS MUST reject an attempt to configure an L2VPN Selected Port with a VLAN ID already assigned by an L2VPN NSI Encapsulation Subtype encoding.

A Point-to-Point forwarding CMTS MUST reject—with a reject-multipoint-L2VPN confirmation code—a registration or service flow transaction that attempts to configure more than one cable attachment circuit (i.e., CM) with the same L2VPN NSI Encapsulation service multiplexing value.

6.2.6 Network System Interface (NSI) Encapsulation

Modern LAN switches and routers implement a rich set of layer 2 bridging features, and wide-area L2VPN operation over MPLS and IP tunneled backbone networks is an active area of product innovation and standardization efforts. This specification does *not* fully specify the layer 2 forwarding of Ethernet packets between CMTSs. It does attempt to specify the configuration of L2VPN forwarding within a single CMTS, and in particular between an RF Interface attachment circuit of a CM and an NSI interface. CMTS vendors are encouraged to support existing and future layer 2 backbone bridging protocols and features when forwarding layer 2 traffic, to and from a DOCSIS RFI MAC interface.

6.2.6.1 NSI Encapsulation Subtype

Although this specification primarily specifies L2VPN operation on the DOCSIS RF interface, it also specifies a limited degree of operation on an NSI interface for the following reasons:

- To standardize the L2VPN configuration for certification testing; and
- To standardize among CMTS vendors a useful subset of L2VPN capabilities.

This specification defines an NSI Encapsulation Subtype of an L2VPN Encoding, Section B.3.2, to optionally describe how an L2VPN's packets are encapsulated on a single selected NSI port. The CMTS vendor implementation may permit this Selected NSI Port to change in the event of port failure or other events. A CMTS vendor MAY use the NSI Encapsulation subtype for additional scenarios, and MAY use vendor specific subtypes of the NSI Encapsulation to support vendor-specific mapping of attachment circuits to backbone pseudo-wires or internal virtual switch instances.

This specification requires CMTSs to implement only a single L2VPN NSI Encapsulation format: [802.1Q] tagging, with a statically configured 12-bit VLAN ID value as the Service Multiplexing value. If the CMTS implements other L2VPN encapsulation formats on an NSI port, it should use the NSI Encapsulation Subtype encoding if the particular format code is defined for the subtype.

6.2.6.2 IEEE 802.1Q L2VPN Forwarding²⁶

When an NSI Encapsulation VLAN ID is statically configured, it is expected to apply to only the Selected Ethernet port. The selection of a particular NSI interface for forwarding a particular L2VPN, or attachment circuit, is vendor-specific. The Vendor-Specific L2VPN Subtype may be used for this purpose.

A point-to-point forwarding CMTS MUST reject a CM registration or service flow transaction with an L2VPN Encoding that omits the NSI Encapsulation subtype or a vendor-specific subtype that identifies the NSI service multiplexing value. A multipoint forwarding CMTS does not require an NSI Encapsulation subtype in an L2VPN Encoding, but MUST accept and implement the subtype if it is specified. A CMTS in either forwarding mode MUST reject a CM registration or service flow transaction with an L2VPN Encoding that contains an NSI Encapsulation subtype for a VPNID that differs from the NSI Encapsulation subtype for that VPNID within any other accepted L2VPN Encoding.

Within the IEEE 802.1Q NSI Encapsulation, VLAN ID values 0, 1, and 4095 are not permitted as a configured VLAN ID. VLAN ID 0 is reserved for priority-only tags in IEEE 802.1Q. VLAN ID 1 is reserved as the default port-based VLAN ID in IEEE 802.1Q. Allowing subscriber L2VPNs to configure onto VLAN ID 1 risks inadvertent layer 2 forwarding of out-of-band management traffic on that subscriber L2VPN. VLAN ID 4095 (all '1's) is reserved by the IEEE.

It is a matter of vendor specific implementation as to whether the CMTS accepts non-L2VPN traffic on an NSI port with a priority-only IEEE 802.1Q tag (i.e., with VLAN ID 0).

6.2.6.3 IEEE 802.1Q L2VPN Forwarding (formerly known as 802.1ad)²⁷

[802.1Q] describes a dual tagging approach (i.e., Q-in-Q) for L2VPN forwarding in a backbone. A packet has an outer 12-bit VLAN ID tag and an inner 12-bit VLAN ID tag. The NSI Encapsulation Subtype allows the pair of 12-bit VLAN ID tags to be configured for each CM or SF performing L2VPN forwarding. The configuration of the dual tags depends on the L2VPN Forwarding Mode of the CMTS and the IEEE 802.1Q networking elements in the backbone.

6.2.6.3.1 Point-to-Point CMTS Forwarding with Point-to-Point 802.1Q Forwarding

In this scenario, the IEEE 802.1Q networking elements in the backbone simply forward point-to-point without learning MAC addresses. The outer 802.1Q tag identifies a destination network element that performs the L2VPN Bridging functions of MAC address learning on a bridge port and forwarding/flooding among those bridge ports. The inner 802.1Q tag identifies a particular bridge port on that external L2VPN Bridge element.

The CMTS is not otherwise configured with the IP address or identity of the destination node in this case. It is configured with only the two VLAN ID tags to use for NSI port encapsulation.

When dual-tagged 802.1Q frames are forwarded in the backbone network, the intermediate nodes only use the outer VLAN ID tag to make forwarding decisions. For example, 802.1Q supports provision of forwarding frames based on the outer tag value without a MAC address lookup.

The L2VPN Bridge addressed by the outer VLAN ID tag is separately configured as to which logical customer L2VPN each inner-tag bridge port is connected.

6.2.6.3.2 Point-to-Point CMTS Forwarding with L2VPN Bridging Network Element

IEEE 802.1Q can be leveraged to construct a backbone networking element to perform the L2VPN switching function of MAC layer learning and forwarding/flooding between attachment circuits belonging to the same L2VPN.

In this scenario, the inner VLAN ID tag represents the logical L2VPN, and the outer VLAN ID tag represents an individual attachment circuit to that logical L2VPN. The IEEE 802.1Q L2VPN switch considers the outer VLAN ID tag to represent a separate virtual trunk interface, while the inner tag represents a logical switch. The 802.1Q L2VPN switch builds an L2 Forwarding Database based on the MAC addresses it learns from each virtual trunk interface,

²⁶ Section heading added per L2VPN-N-12.1065-11 on 9/5/12 by JB.

²⁷ Section and subsections modified per L2VPN-N-14.1186-1 on 1/15/15 by PO.

and forwards/floods packets among those virtual trunk interfaces. In this manner, it provides L2VPN switch forwarding between all attachment circuits of an L2VPN. Using this technique, over 4000 L2VPN service instances can be supported between one CMTS and the IEEE 802.1Q L2VPN Switch, and each of these can have over 4000 cable modems/service flow associations at the CMTS.

6.2.6.4 IEEE 802.1Q L2VPN Forwarding (formerly known as 802.1ah)²⁸

[802.1Q] describes Provider Backbone Bridging (PBB) (i.e., MAC-in-MAC) which introduces a hierarchical sub-layer, by encapsulating the customer Ethernet frame within a service provider frame. PBB backbone edge switches append their own source address (B-SA) and destination address (B-DA), as well as a backbone VID (B-VID). A service ID (I-SID) field identifies a customer-specific service instance.

The NSI Encapsulation Subtype allows these values to be configured for each CM or SF performing L2VPN forwarding. The configuration of the tags depends on the L2VPN Forwarding Mode of the CMTS and the IEEE 802.1Q networking elements in the backbone.

6.2.7 Virtual Private LAN Service (VPLS) and Virtual Private Wire Service²⁹

The IETF has defined L2VPN framework [RFC 4664] by two types - multipoint L2VPN (VPLS) and point-to-point L2VPN (VPWS), using the notion of a Pseudowire (PW) [RFC 3985] to implement L2VPN over a Packet Switched Network (PSN) e.g., IP or MPLS network. The interoperation of DOCSIS L2VPN Forwarding with PW is required to allow seamless L2VPN forwarding over an MSO's IP/MPLS network.

Following the PW architecture [RFC 3985], the CMTS acts as a Provider Edge (PE) router having L2VPN Forwarder function, and does a simple mapping between the PW and Attachment Circuit (AC) based on local information, i.e., PW demultiplexer and incoming/outgoing logical/physical port. From the DOCSIS network standpoint, each CM with at least one upstream forwarding SF for an L2VPN corresponds to an AC at the CMTS. As described in Section 5.2.2, a cable AC to an L2VPN is considered to be the CM, not an individual service flow on the CM. Ethernet frames received over the L2VPN attachment circuits using either the MPLS or L2TPv3 NSI Encapsulation subtype are intended to be forwarded by the CMTS on PWs.

A PW emulates the essential attributes of an L2VPN service (e.g., Ethernet) over a packet switched network (PSN) (e.g., IP or MPLS network). Simply put, a PW is a tunnel with either MPLS encapsulation or L2TPv3 encapsulation, which is used to carry the service-specific PDUs (e.g., Ethernet frames) through the PSN. An L2VPN instantiation means establishment of one or more pairs of PWs between two or more L2VPN Forwarders (i.e., PE routers) in the PSN.

In VPWS, an L2VPN Forwarder (i.e., PE router) binds a PW to a single AC, such that frames received on the one are sent on the other, and vice versa. VPWS involves two PE routers logically connected via a bidirectional PW.

In VPLS, an L2VPN Forwarder (i.e., PE router) binds a set of PWs to a set of ACs; when a frame is received from any member of that set, a MAC (Media Access Control) address table is consulted (and various 802.1d procedures executed) to determine the member or members of that set on which the frame is to be transmitted. VPLS involves two or more PE routers logically connected via distinct PWs.

PW can be set up on PE routers using one of two methods: (1) Static Configuration or (2) Dynamic Signaling, enforcing one of two data plane encapsulations. This is covered in detail in Section 6.2.7.2.

A PSN tunnel needs to be set up on PE routers for PWs to be useful. PSN tunnel data plane and control plane are discussed in Section 6.2.7.1.

6.2.7.1 PSN Tunnel – Control Plane and Data Plane Encapsulation

In order for Layer 2 VPN services to work in a PSN, a functioning IP or MPLS data plane between the PE routers having L2VPN Forwarder functions needs to exist. After all, the PW tunnels travel through PSN tunnels. The PSN tunnel may be an MPLS tunnel (i.e., MPLS Encapsulation) or IP tunnel (e.g., L2TPv3 Encapsulation or other encapsulation not defined in this specification). This means that a CMTS acting as a PE router in the PSN needs to support the corresponding tunnel data plane and control plane.

²⁸ Section added per L2VPN-N-12.1065-11 on 9/5/12 by JB, L2VPN-N-14.1186-1 on 1/15/15 by PO.

²⁹ Section revised per L2VPN-N-12.1065-11 on 9/5/12 by JB.

A CMTS MUST support MPLS data plane, as defined in [RFC 3031] and [RFC 3032]. That is, a CMTS MUST support MPLS packet forwarding (e.g., receiving and transmitting MPLS packets having ethertype=0x8847 and 0x8848) on its NSI ports. An MPLS packet is assumed to have one or more label stack entries.

A CMTS MUST support LDP as the MPLS control plane, as defined in [RFC 5036] and [ID LDPv6], on its NSI port. This is required to setup MPLS Label Switched Paths (LSPs) through dynamic label allocation and label exchange (for IGP prefixes) with other routers in PSN.

A CMTS SHOULD support BGP [RFC 3107] as the MPLS control plane to set up MPLS LSPs (for BGP prefixes identifying PE routers) with other routers in PSN. BGP is beneficial for higher scale.

A CMTS MAY support RSVP-TE [RFC 3209] as the MPLS control plane to setup MPLS LSPs for IGP prefixes. RSVP-TE is beneficial to construct path-constrained LSPs.

A CMTS MUST provide a mechanism to limit the advertisement of LDP Forward Equivalence Class of only host entries.

A CMTS SHOULD support TCP MD5 authenticity and integrity based on the use of the TCP MD5 Signature Option specified in [RFC 5925] per LDP neighbor.

A CMTS MAY support LDP Fast Reroute (FRR) as defined in a Basic Specification for IP Fast Reroute: Loop-Free Alternates [RFC 5286].

The usage of any other PSN Tunnel (i.e., IP Tunnel) not defined in this specification is not recommended for L2VPN packet forwarding.

6.2.7.2 PW – Control Plane and Data Plane Encapsulations

A PW needs to be set up between two PE routers participating in L2VPN. A PW can be established using one of two methods: (1) Static Provisioning or (2) Dynamic Signaling. The latter method is known as having a PW control plane.

The static method requires no protocol and relies on static assignment of PW encapsulation values (i.e., MPLS labels or L2TPv3 session ID) on CMTSs using either vendor-specific CMTS configuration or L2VPN vendor-specific Subtype parameters. Please note that the static method is not recommended.

The dynamic signaling method relies on the CMTS using a protocol to signal PW encapsulation values (i.e., MPLS label or L2TPv3 session ID) so as to set up PW with remote PE router(s). The following three signaling protocols are standardized by the IETF for PW establishment:

- LDP [RFC 4447] [RFC 4762]
- L2TPv3 [RFC 3931] [RFC 4667]
- BGP [RFC 4761] [RFC 6624] (See Note below)

NOTE: Usage of BGP as a signaling protocol is tied into the usage of BGP as the auto-discovery protocol.

Each of the above three signaling protocols require the CMTS (i.e., PE router) to have the following information per PW per L2VPN for successful PW setup with a remote PE:

- 1) Remote PE Identifier (e.g., IP address of each remote PE)
- 2) PW Identifier
- 3) Other Parameters (e.g., MTU)

This information can be either encoded using L2VPN encodings, or dynamically formulated by the CMTS based on auto-discovery mechanism (e.g., BGP). The next two sub-sections focus on specifying the usage of PW signaling protocols with or without auto-discovery.

The CMTS SHOULD support the PW Control Word for use within a MPLS Network as defined in [RFC 4385] and [RFC 4448].

6.2.7.2.1 *PW Signaling without Auto-Discovery*³⁰

If the NSI Encapsulation subtype of MPLS or L2TPv3 is included for any L2VPN and the BGP Attribute sub TLV (43.5.21) is NOT included, the CMTS MUST NOT initiate the usage of any auto-discovery procedures for that L2VPN, irrespective of the auto-discovery protocol enabled on the CMTS.

When the NSI Encapsulation subtype is present, the CMTS uses the encapsulation protocol (e.g., MPLS or L2TPv3) that is specified on a per L2VPN basis in the data plane, along with the corresponding signaling protocol (e.g., LDP or L2TPv3) in the control plane to set up the PW with the remote PE router whose IP address is specified in NSI Encapsulation subtype field.

A CMTS MUST support NSI Encapsulation subtype of MPLS. The CMTS MAY support NSI Encapsulation subtype of L2TPv3.

Using MPLS encapsulation implicitly means that a CMTS MUST support MPLS encapsulation of Ethernet Frames, as per [RFC 4448]. The CMTS would encapsulate the upstream Ethernet frames received on cable attachment circuit (e.g., SF) into an Ethernet over MPLS PW [RFC 4448] and forward them via the NSI port. Similarly, the CMTS decapsulates the downstream MPLS packets received on NSI port and forwards the resulting Ethernet frames to the CM via the downstream service flow (e.g., SF).

Using MPLS encapsulation implicitly means that a CMTS MUST support LDP for PW signaling, as per [RFC 4447] and [RFC 4762]. Specifically, the CMTS MUST use PwID FEC Element to establish the PW using LDP. The value of PwID is learned from the Pseudowire ID Encoding. The CMTS SHOULD NOT use Generalized PwID FEC Element (e.g., SAIL in B.3.6 and TAIL in B.3.7) to establish the PW using LDP, as it is meant to be used with auto-discovery by CMTS.

A CMTS MAY support L2TPv3 for PW signaling and L2TPv3 Encapsulation, as per [RFC 3931].

If LDP is used for PW establishment, the PSN tunnel is to be an MPLS tunnel (e.g., an MPLS data plane on the NSI of CMTS).

6.2.7.2.2 *PW Signaling with Auto-Discovery*³¹

Auto-discovery of PE routers participating in a given L2VPN is done using Multi-Protocol BGP extensions for the L2VPN address-family.

If the NSI Encapsulation subtype and the BGP Attribute sub TLV (43.5.21) are included for an L2VPN, then the CMTS MUST initiate MP-BGP based auto-discovery procedures for that L2VPN. The MPLS PW ID, MPLS Peer IP address and L2TPv3 TLV values will be ignored by the CMTS in this case. If the NSI Encapsulation is not included the CMTS will assume MPLS NSI Encapsulation and will initiate MP-BGP auto-discovery procedures accordingly.

The CMTS MAY use vendor-specific configuration to dynamically select and learn the Service Multiplexing value for an L2VPN Encoding from the CMTS's L2VPN peers.

The CMTS could signal the PW setup using LDP or L2TPv3 immediately after the auto-discovery is completed (using BGP), as specified in [RFC 6074], or signal the PW setup using BGP at the same time as that of auto-discovery (using BGP), as specified in [RFC 4761] and [RFC 6624]. While both approaches could be utilized, they result in using either two protocols (BGP and LDP) or one protocol (BGP).

A CMTS MUST support BGP for auto-discovery and LDP for PW signaling, as per [RFC 6074] [RFC 4447].

A CMTS MAY support BGP for PW signaling and auto-discovery, as per [RFC 4761] and [RFC 6624], by implementing the MP-BGP L2VPN address-family. Specifically, a CMTS MAY support [RFC 4761] for Multipoint L2VPN and [RFC 6624] for Point-to-Point L2VPN.

A CMTS MAY support BGP for auto-discovery and L2TPv3 for PW signaling, as per [RFC 6074] and [RFC 4667].

³⁰ Revised per L2VPN-N-13.1109-3 on 7/22/13 by PO.

³¹ Revised per L2VPN-N-13.1109-3 on 7/22/13 by PO.

When implementing BGP, the CMTS MAY implement BGP Support for Four-octet AS Number Space [RFC 4893]. The CMTS MUST support the co-existence of two-octet and four-octet BGP autonomous system numbers (ASN). The CMTS MUST support Route Refresh Capability for BGP-4 [RFC 2918].

The CMTS MAY support the TCP Authentication Option for BGP [RFC 5925].

When implementing BGP auto-discovery, the CMTS treats the L2VPN as a Point-to-Point L2VPN, unless it receives the Virtual Switch Instance encoding (Section B.3.20), in which case it treats the L2VPN as a Multipoint L2VPN.

The CMTS MUST deploy one of the two PW signaling with auto-discovery methods as follows:

- If the BGP VPN ID TLV (43.5.21.1) is included then LDP PW signaling with BGP auto-discovery [RFC 6074] will be deployed.
- If the BGP CE-ID/VE-ID TLV (43.5.21.5) is included then auto-discovery and PW signaling with BGP [RFC 4761] and [RFC 6624] will be deployed.

Once the CMTS has instantiated a particular L2VPN, the CMTS invokes BGP to advertise the availability of the L2VPN site, as per [RFC 6074] or [RFC 4761] and [RFC 6624].

When BGP auto-discovery with LDP signaling is used [RFC 6074], the CMTS auto-generates Route Distinguisher (RD), Route-Target (RT), L2VPN Extended Community and L2VPN NLRI that are encoded in a BGP UPDATE message using the 4-byte BGP VPNID (Section B.3.21.1) and using local information such as BGP Router-ID, as described below.

- 1) The L2VPN Extended Community value is calculated by combining 2-byte BGP ASN with 4-byte BGP VPNID in this format ASN:BGP VPNID.
Note that the BGP VPNID is assumed to be an AS-wide unique number.
- 2) The RD value is the same as that of L2VPN Extended Community.
- 3) The RT value is the same as that of L2VPN Extended Community (and used as both import RT and export RT values), unless CM conveys import and export RT to CMTS via TLV 43.5.21.3 and 43.5.21.4.
- 4) The L2VPN NLRI is calculated by combining the RD (calculated above) with a 4-byte number formatted as follows: RD: POOL_NUMBER, where the 4-byte CMTS BGP Router-ID is used as the POOL_NUMBER.

According to [RFC 4761] and [RFC 6624], when BGP auto-discovery and PW signaling is used for a L2VPN, the CMTS will encode in a BGP UPDATE message a specific L2VPN/VPLS NLRI, Route-Targets and an additional Layer 2 Info Extended Community, as described below:

- 1) The L2VPN NLRI will be generated by the CMTS as follows:
 - a. The RD value is the same as that of the Route Distinguisher TLV (43.5.21.2).
 - b. The VE ID value is the same as that of the BGP CE-ID/VE-ID TLV (43.5.21.5). Please note that if the L2VPN instance is a Virtual Switch Instance and multiple local CMs are part of this VSI then the NLRI might have multiple VE-IDs/CE-IDs.
 - c. The rest of the L2VPN/VPLS NLRI fields (VE Block Offset, VE Block Size and Label Base) will be generated by the CMTS based on internal variables.
- 2) The RT value is the same as that of the RD (and used as both import RT and export RT values), unless CM conveys import and export RT to CMTS via TLV 43.5.21.3 and 43.5.21.4.
- 3) The L2 Info Extended Community will be generated by the CMTS as follows:
 - a. Encapsulation Type will be set to point-to-point Ethernet RAW Mode (5), unless a different value is conveyed by the CM via the MPLS PW Type TLV (43.5.2.4.3). Acceptable values: Ethernet Tagged Mode (4), RAW Mode (5) and VPLS(19),
 - b. The rest of the fields (Control Flags and Layer2 MTU) will be generated by the CMTS based on internal variables.

Following [RFC 6074] or [RFC 4761]/[RFC 6624] upon receiving a BGP UPDATE message from any of the remote PE routers, the CMTS checks for the matching RT value (e.g., checks the received RT value with the local RT value that was auto-generated). If a match is found, then the CMTS concludes to have discovered the remote PE router for that L2VPN. The CMTS invokes LDP to setup the PW with that remote PE router (identified by its IP address that was in BGP NEXT_HOP attribute), as per [RFC 6074] or, in the [RFC 4761]/[RFC 6624] case, starts using the information in the L2VPN/VPLS NLRI and the L2 Info Extended Community to setup the PW.

If CMTS instantiated a P2P L2VPN, then CMTS SHOULD setup the PW with only one remote PE using LDP signaling [RFC 4447] or BGP signaling [RFC 6624]. If CMTS instantiated a multipoint L2VPN, then CMTS SHOULD setup the PW with one or more remote PEs using LDP signaling [RFC 4447] or BGP signaling [RFC 4761], as/when they get discovered using MP-BGP.

If the CMTS discovers more than one remote PE router for an L2VPN that is locally marked P2P L2VPN, then the CMTS MUST tear down any pre-established PW for that L2VPN. In such case, the CMTS SHOULD generate an error (e.g., syslog). This helps to avoid any misconfiguration of remote PE routers and avoids possible L2VPN leakage. Note that there is nothing in BGP that identifies a given advertisement as P2P or Multipoint L2VPN.

A CMTS MUST support LDP for PW signaling, as per [RFC 4447]. Specifically, the CMTS MUST use Generalized PWid FEC Element (e.g., AGI, SAI, TAI) to establish the PW using LDP, as per [RFC 6074]. This means that AGI = L2VPN Extended Community Value, SAI = Local POOL_NUMBER and TAI = Remote POOL_NUMBER.

A CMTS MUST support MPLS Encapsulation, as per [RFC 4448]. In data plane, the CMTS encapsulates the upstream Ethernet frames received on an L2VPN Service Flow into an Ethernet over MPLS PW [RFC 4448]. Similarly, the CMTS decapsulates the downstream MPLS packets received on NSI port and forwards the resulting Ethernet frames to the CM via the downstream service flow.

If LDP is used for PW establishment, the PSN tunnel is an MPLS tunnel (e.g., an MPLS data plane on the NSI of CMTS).

6.3 CMTS Upstream L2VPN Forwarding ³²

The CMTS MUST NOT interpret an 802.1Q tag already appearing in an upstream packet as providing either the priority or L2VPN identifier for L2VPN forwarding bridging. This includes a priority-only tag. The CMTS MUST transparently forward any subscriber-provided 802.1Q tag independent of the NSI encapsulation. If the subscriber-tagged packet is forwarded with 802.1Q NSI Encapsulation on an 802.1Q NSI Ethernet port, the CMTS MUST prepend an outer 802.1Q tag before the inner subscriber-provided tag.

A CMTS MUST be able to send and receive for L2VPN forwarding on all interfaces packets as large as the maximum supported packet PDU length as defined in [DOCSIS MAC]. The maximum packet PDU length supported in a pre-3.1 DOCSIS CMTS is 1522 bytes and in a DOCSIS 3.1 CMTS is 2000 bytes.

For example, on an Ethernet NSI interface accepting 802.1Q tags for L2VPN NSI Encapsulation, the CMTS accepts and forwards a 1526 byte (pre-3.1 DOCSIS) or a 2004 byte (DOCSIS 3.1) Ethernet packet. For pre-3.1 DOCSIS, such a packet is forwarded to the downstream RF MAC domain as a 1522 byte packet consisting of a nominal maximum length 1518 Ethernet packet with one four-byte subscriber non-service delimiting tag. For DOCSIS 3.1 such a packet is forwarded to the downstream RF MAC domain as a 2000 byte packet consisting of a nominal maximum length 1996 Ethernet packet with one four-byte subscriber non-service delimiting tag.

The CMTS MUST NOT count learned L2VPN CPE MAC addresses learned from upstream packets towards any enforced docsSubMgtCpeControlMaxCPEIp setting for the CM [DOCSIS MAC]. This setting applies only to learned subscriber IP addresses that are forwarded in a non-L2VPN manner. Multipoint mode CMTSs have a separate requirement to limit the number of source MAC addresses learned on each L2VPN.

The CMTS MUST NOT apply subscriber management filtering [DOCSIS MAC] to upstream L2VPN forwarded packets.

³² Section modified per L2VPN-N-07.0471-2 on 7/5/07 by KN and further revised per L2VPN-N-12.1065-11 on 9/19/12 by CP. Revised per L2VPN-N-14.1135-5 on 3/27/14 by JB.

The CMTS MUST NOT perform the TOS Overwrite function [DOCSIS MAC] for upstream L2VPN-forwarded packets.

An L2VPN forwarder maintains an out-of-band 3-bit user priority, as defined in Section 6.7.2.2, associated with each forwarded packet. The CMTS MUST encode the egress value of user priority as specified for the NSI Encapsulation format (e.g., in the user-priority bits of an IEEE 802.1Q tag or in the Traffic Class bits of an MPLS header). The CMTS SHOULD provide mapping of upstream user priority to NSI port transmission traffic class as specified for the NSI Encapsulation format. The number of NSI port transmission traffic classes is vendor-specific. If an upstream service flow L2VPN Encoding omits the Upstream User Priority subtype, the CMTS MUST by default forward such packets to an NSI port with a user priority of zero. The CMTS MAY forward with non-zero default user priority values with vendor-specific configuration. The CMTS MAY also use this out-of-band 3-bit user priority of received NSI encapsulated packets for downstream classification (e.g., in the case of multipoint L2VPN forwarding).

The CMTS MUST support both L2VPN and non-L2VPN forwarding of upstream traffic from a Forwarding L2VPN service flow based on checking the source MAC address against a CM Interface Mask configured for the SF. The CMTS MUST direct packets from the source MAC addresses indicated with a '1' in the CM Interface Mask to the L2VPN forwarder. The CMTS MUST NOT direct to the L2VPN Forwarder packets from the eCM and at least one other eSAFE source MAC address, even when received on an L2VPN forwarding upstream service flow, when the interfaces corresponding to those source MAC addresses are indicated with a '0' in the CM Interface Mask.

A CMTS MAY recognize when the CM Interface Masks in the set of Upstream Packet Classifier L2VPN Encodings of a compliant CM permit it to avoid checking upstream source MAC addresses and instead forward upstream packets to the L2VPN or non-L2VPN forwarder solely on the basis of the upstream service flow.

The CMTS MUST recognize a CM Interface Mask criterion in a Downstream Packet Classifier L2VPN Encoding regardless of whether the Encoding classifies L2VPN or non-L2VPN traffic. The CMTS MUST classify the destination MAC address of a downstream packet as one of three classes: 1) A CM MAC address; 2) a CPE MAC address; or 3) an eSAFE MAC address of a particular eSAFE host type. The CMTS MUST consider the criterion to be matched when the destination MAC address of the downstream layer 2 packet is a CM MAC address or eSAFE MAC address that corresponds to a host type with a '1' in the CM Interface Mask. The CMTS MUST consider the criterion to be matched when the destination MAC address is a CPE MAC address and the CM Interface Mask has *any* CPE host type bit set, i.e., any of bits 1 or 5-15 set. The CMTS MUST consider the criterion to be unmatched when the destination MAC address is a CM or eSAFE MAC address and the single CM Interface mask bit corresponding to that host type has a '0' bit. The CMTS MUST consider the criterion to be unmatched when the destination MAC address is a CPE MAC address and the CM Interface Mask has a zero bit in all CPE host type positions, i.e., has a zero bit in positions 1 and 5-15.

The CMTS MUST reject any attempt (i.e., registration or DSx transaction) to configure multiple Upstream Classifier L2VPN Encodings that classifies to the same upstream Service Flow but with a different VPNID Subtypes. The CMTS uses the upstream SF to determine a single VPNID for L2VPN forwarding.

6.4 CMTS Downstream L2VPN Forwarding³³

The CMTS MUST reject any REG-REQ or REG-REQ-MP with an L2VPN Encoding if BPI is not also enabled in the REG-REQ or REG-REQ-MP. The CMTS MUST reject any DSA-REQ or DSC-REQ with an L2VPN Encoding if BPI is not also enabled for the CM.

The CMTS MUST NOT apply subscriber management filters ([DOCSIS MAC]) to downstream L2VPN forwarded traffic.

³³ Section modified per L2VPN-N-0741-2 on 7/5/07 by KN and revised per L2VPN-N-12.1065-11 on 9/19/12 by CP.

The CMTS MUST accept a single Downstream Classifier L2VPN Encoding in a Downstream Packet Classification Encoding of a REG-REQ, REG-REQ-MP, DSA-REQ, or DSC-REQ message. A CMTS MUST apply classifier rules that contain an L2VPN Encoding only to packets forwarded by the L2VPN Forwarder. Furthermore, only classifiers containing an L2VPN Encoding may be applied against downstream L2VPN forwarded traffic.

- The CMTS MUST reject the registration of cable modems with invalid Downstream Packet Classification Encodings.
- A valid Downstream Classification L2VPN Encoding contains zero or one VPNID subtypes, zero or one User Priority Range subtypes, and any number of Vendor Specific L2VPN Parameter subtypes. The CMTS MUST silently ignore all invalid L2VPN Encoding subtypes.
- The CMTS MUST accept as valid and silently ignore any unrecognized L2VPN Encoding subtypes.
- The CMTS MUST accept multiple Downstream Classification Configuration Settings with a Downstream Classifier L2VPN Encoding that classify different L2VPN VPN IDs to the same referenced service flow.
- The CMTS MUST support the same Classifier criteria options for Downstream Classifier L2VPN Encodings as it does for non-L2VPN Downstream Classifier L2VPN Encodings.
- The CMTS MUST interpret a Downstream Packet Classifier Encoding containing no other criteria than a Classifier L2VPN encoding as matching *all* packets forwarded downstream on the L2VPN identified by the VPNID of the Classifier L2VPN Encoding, and classify all such packets to the referenced service flow.
- The CMTS MAY accept multiple Downstream Classifier L2VPN Encodings with the same VPNID classifying packets to different service flows. Operation is undefined when more than one Downstream Classifier matches a particular downstream packet.
- The CMTS MUST reject a service flow transaction request containing an invalid Downstream Classifier L2VPN Encoding.
- If the L2VPN Encoding contains a Downstream User Priority Range subtype, the CMTS MUST match the classifier only to L2VPN forwarded packets that are received on NSI with the user priority within the indicated range. Otherwise, the classifier applies to all user priorities.

With the acceptance of a valid Downstream Classifier L2VPN Encoding, the L2VPN forwarder of the CMTS MUST forward on the referenced service flow of the classifier all single-CM downstream traffic destined for CPE attached to that CM. For Point-to-Point mode, this means all downstream traffic on the L2VPN; for Multipoint mode this means unicast traffic destined to CPE MAC addresses learned from upstream traffic from the CM. If downstream L2VPN forwarded traffic is not classified to a particular downstream service flow, the CMTS MUST forward single-CM traffic on the CM's primary downstream service flow.

The CMTS MUST be able to classify layer 2 packets as they would appear on the RFI interface; that is, NOT including any service-delimiting encapsulation header (e.g., 802.1Q tag, MPLS, IP) that appeared on the CMTS NSI port. This means that the 802.1Q Packet Classification Encodings [DOCSIS MAC] apply only to the subscriber-private or inner 802.1Q tag, and not the NSI encapsulation header of the L2VPN-forwarded packet. On the other hand, the Downstream User Priority Range subtype that may be included in the downstream Service Flow Packet Classifier encoding applies only to the service delimiting encapsulation header (e.g., 802.1Q, MPLS, IP).

A CMTS MUST forward downstream L2VPN-forwarded packets for different L2VPNs on different downstream service flows. This provides isolation of QOS for L2VPN service.

Unless explicitly configured to combine the forwarding data base of different L2VPNs, the CMTS L2VPN Forwarder MUST maintain upstream and downstream separation of L2 forwarded traffic between attachment circuits configured with different VPNIDs. The number of CMs or SFs supported for L2VPN Forwarding is CMTS vendor-specific. The number of unique VPNIDs supported by a CMTS is vendor-specific.

6.4.1 Multipoint Downstream Forwarding

The CMTS MUST reject (with a reject-permanent confirmation code) a registration or service flow transaction that would require defining L2VPN SAIDs exceeding the CM's Downstream SAID capability ([DOCSIS MAC]).³⁴

A Multipoint forwarding mode CMTS MUST learn the source MAC addresses of upstream CPE traffic and associate them with a particular CM on that L2VPN.

A Multipoint forwarding CMTS MUST limit the number of MAC addresses permitted to be learned on any single L2VPN to a configurable value that applies to all L2VPNs. The CMTS SHOULD permit configuration of the maximum number of MAC addresses per L2VPN on a per-L2VPN basis.

A Multipoint forwarding CMTS MUST forward downstream packets encrypted on different L2VPN SAIDs on different service flows.

6.4.2 DOCSIS 3.0 L2VPN Downstream Multicast Forwarding³⁵

The downstream forwarding of unicast traffic on DOCSIS 3.0 devices is the same as earlier versions of DOCSIS. DOCSIS 3.0 introduces the concept of Multicast DSID Forwarding for downstream multicast traffic [MULPIv3.0]. When Multicast DSID Forwarding is enabled, the CMTS is required to label all multicast traffic with a DSID, which it communicates to the CM. The CMTS communicates the resequencing and multicast attributes of the DSID to the CM in the DSID encodings (TLV 50) of the Registration Response message or the Dynamic Bonding Change message. The CM filters downstream multicast traffic which is labeled with an unknown DSID or is not labeled with a DSID.

When Multicast DSID Forwarding is enabled on a CM, the CMTS MUST label all L2VPN multicast traffic intended for that CM with a DSID. The CMTS could meet this requirement in different ways. For example,

- the CMTS could sequence all L2VPN traffic (as the Resequencing DSID header is present in all sequenced traffic); or
- the CMTS could label the L2VPN multicast traffic with an "L2VPN DSID" which is assigned per CM, per L2VPN instance, or even globally.

This specification introduces the concept of an "L2VPN DSID" to enable the CMTS to label the downstream multicast L2VPN-forwarded traffic that is not otherwise sequenced for a bonded service flow. If L2VPN DSIDs are used, the CMTS MUST communicate L2VPN DSIDs to the CM in the L2VPN Encodings (TLV 43.5) of the Registration Response message or the Dynamic Service message. If L2VPN DSIDs are used, the CMTS MUST NOT communicate an L2VPN DSID to the CM in the DSID encodings (TLV 50) of the Registration Response or Dynamic Bonding Change message.

6.5 L2VPN Isolation and Privacy

A key goal of this specification is to *isolate* traffic between L2VPN and non-L2VPN subscribers, as well as between different L2VPN subscribers. Non-L2VPN (i.e., residential) subscribers should not be able to see traffic forwarded to L2VPN subscribers, and L2VPN subscribers, moreover, should not see traffic intended for non-L2VPN residential subscribers.

6.5.1 Protecting L2VPN Traffic

This specification uses BPI encryption to isolate L2VPN from non-L2VPN traffic in the downstream. This requires a cable operator to configure CMs providing L2VPN service to enable Baseline Privacy Interface (BPI) operation. The cable operator is expected to configure all CMs, with and without L2VPN forwarding, to enable BPI, so that all such CMs can receive encrypted IP multicast traffic.

The CMTS MUST reject an attempt (i.e., a registration or DSx transaction) to configure a forwarding L2VPN Encoding if the CM is not also configured to support BPI operation.

³⁴ Section modified per L2VPN-N-07.0471-2 on 7/5/07 by KN.

³⁵ Section added per L2VPN-N-12.1065-11 on 9/19/12 by CP.

A CMTS MUST assign at least one L2VPN SAID for downstream forwarding to each separate L2VPN forwarded by the CMTS on a downstream channel. A single L2VPN SAID assigned for all CMs on the same L2VPN is called a Group L2VPN SAID. The CMTS MAY assign L2VPN SAID values to be different for the same L2VPN on different downstream channels. A CMTS MAY assign multiple L2VPN SAIDs to the same L2VPN on the same downstream channel, e.g., to assign an Individual L2VPN SAID to each CM in Point-to-Point forwarding mode. The CMTS MUST assign a Group or Individual L2VPN SAID that differs from any other Primary SAID assigned on that channel. A CMTS MAY assign multiple SAIDs to the same L2VPN on the same CM.

A CMTS MUST add to the forwarding L2VPN Encoding of its REG-RSP or REG-RSP-MP and Dynamic Service messages to an L2VPN-compliant CM one or more L2VPN SA-Descriptor Subtypes for its assigned L2VPN SAID(s) for downstream forwarding to that L2VPN on the CM's downstream channel. The CMTS MUST encode separate top-level L2VPN encodings for each separate L2VPN ID. A CMTS MAY add L2VPN SA-Descriptor Subtypes in messages to non-compliant CMs, but they will be ignored by the CM. The CMTS MUST describe the L2VPN SAIDs with an SA-Type of Dynamic in an L2VPN SA-Descriptor Encoding.³⁶

A CMTS does NOT include SA Descriptors for all L2VPN SAIDs it assigned for a registering modem in its initial BPI Authorization Reply to the CM after registration.

A CMTS MUST encrypt all downstream L2VPN forwarded traffic in an L2VPN SAID assigned to the L2VPN.

The CMTS MUST NOT forward downstream L2VPN traffic to a CM until that CM has completed BPI Authorization and TEK negotiation for the L2VPN SAID in which the traffic is to be encrypted.

A Point-to-Point forwarding CMTS SHOULD assign the same Group L2VPN SAID to different CMs on the same MAC domain attaching to the same L2VPN Identifier, but MAY choose to assign an Individual L2VPN SAID unique for the CM.

A Multipoint forwarding CMTS MUST assign at least one broadcast L2VPN SAID to all CMs on the same MAC Domain attaching to the same L2VPN Identifier. The Multipoint forwarding CMTS MUST forward downstream broadcast packets of the L2VPN encrypted on such a broadcast L2VPN SAID.

A CMTS MAY support vendor-specific configuration to dynamically start or discontinue L2VPN forwarding through a registered CM. A CMTS that discontinues L2VPN forwarding through a CM MUST dynamically delete all upstream service flows forwarding to that L2VPN, and MUST signal a top-level L2VPN encoding to the CM that omits all SA-Descriptors for that L2VPN. This signals the CM to discontinue downstream decryption for the L2VPN SAIDs associated with the L2VPN.

6.5.2 Preventing Leaking of non-L2VPN Traffic

One issue with L2VPN operation is downstream non-L2VPN layer 2 traffic to a Group MAC (GMAC) address, i.e., a layer 2 broadcast or multicast. Downstream non-L2VPN broadcast traffic includes CMTS-originated ARPs to non-L2VPN CPEs and CMTS router advertisements for RIP or OSPF. By default, *all* cable modems—L2VPN and non-L2VPN—will forward to their CPE interface downstream broadcast traffic that is *not* encrypted. Furthermore, unencrypted non-L2VPN GMAC traffic sent to the same multicast Ethernet Destination Address used by a private L2VPN would also be forwarded by an L2VPN CM to its CPE interface. Without special attention, downstream non-L2VPN GMAC traffic will leak onto the L2VPN subscriber's supposedly private CPE network.

This specification addresses the non-L2VPN GMAC leakage problem with the following mechanisms:

- Downstream Unencrypted Traffic (DUT) Filtering
- Downstream IP Multicast Encryption (DIME)

³⁶ Paragraph modified per L2VPN-N-07.0741-2 on 7/5/07 by KN.

6.5.2.1 Downstream Unencrypted Traffic (DUT) Filtering

A cable operator can prevent the leaking of clear-text, non-L2VPN traffic, through L2VPN-compliant CMs by enabling the Downstream Unencrypted Traffic (DUT) Filtering Encoding. When DUT Filtering is enabled, a DUT CM Interface Mask (DUT CMIM) is defined to limit the forwarding of downstream unencrypted traffic to only the interfaces with a '1' bit for that interface in the CMIM. The DUT CMIM by default contains '1' bits only for the internal eCM and eSAFE host interfaces, requiring the CM to prevent forwarding of such traffic to the CMCI interface(s) on the CM. DUT filtering alone prevents non-L2VPN GMAC leaking onto an L2VPN subscriber's CPE network.

6.5.2.2 Downstream IP Multicast Encryption (DIME)

CMs by default has to enable GMAC promiscuous forwarding for L2VPN operation, and in most DOCSIS 2.0 and earlier CMs, this causes all unencrypted GMAC traffic to be delivered to CM software.³⁷ Although the DUT Filtering of the CM's software indeed prevents this traffic from leaking onto the CPE interface, if it is significant, it may affect the forwarding performance of the desired L2VPN traffic through the CM. It is desirable to allow the L2VPN CMs to use their hardware-based SAID filters to drop high-volume non-L2VPN GMAC traffic.

In most CMTS deployments, it is expected that IP multicast session traffic will be the most significant source of high-volume downstream GMAC traffic. It is desirable to encrypt this traffic in a SAID that is unknown by the L2VPN CMs, so that their hardware will filter the packets before delivering it to L2VPN software.

A CMTS MUST implement a configurable option to enable or disable Downstream IP Multicast Encryption (DIME). With DIME enabled, the CMTS MUST encrypt all non-L2VPN downstream IP Multicast traffic that is either statically joined with the BPI+ MIB or dynamically joined with upstream SA-MAP requests from a CM. DIME does not require the CMTS to encrypt non-L2VPN downstream *unjoined* IP multicasts, (e.g., RIPv2 or OSPF multicasts).

By encrypting joined non-L2VPN IP multicast traffic in a non-L2VPN SAID, the potentially high volume downstream non-L2VPN multicast traffic is required to be filtered by DOCSIS 2.0 CMs implementing this specification.

Since the non-L2VPN GMAC traffic is encrypted in a SAID unknown to the L2VPN CM, the CM filters the downstream traffic and prevents it from leaking onto the private CPE network.

6.5.2.3 Mixing L2VPN and non-L2VPN forwarding on the same CM

The L2VPN feature supports mixing of L2VPN and non-L2VPN forwarding for different CPE hosts connected to the same CM. In this case, the L2VPN and non-L2VPN traffic are of course not isolated on the CMCI network(s) of the CM, both in the upstream and in the downstream direction. To support mixed L2VPN and non-L2VPN, the cable operator can configure Upstream L2VPN Classifiers in the CM with a rule that identifies the particular type of traffic to be forwarded on the L2VPN, (e.g., traffic with the source MAC address of a particular CPE).³⁸

This specification does not require the CM to restrict the forwarding between L2VPN and non-L2VPN CPE hosts when they are connected on different CMCI ports of a CM. The embedded VLAN (eVLAN) model of CM forwarding, if implemented on the CM, can provide this isolation (Appendix III).

DUT filtering should not be enabled when mixing L2VPN and non-L2VPN CPE hosts on the same CM, because it is necessary for the non-L2VPN CPEs to still receive downstream ARPs and DHCP broadcasts. With DUT Filtering disabled, however, all downstream non-encrypted GMAC traffic will pass to the mixed CPE network. To prevent this, Downstream IP Multicast Encryption (DIME) can be enabled to prevent the forwarding of unjoined multicast traffic to the mixed-mode CPE LAN.

6.6 CM and eSAFE Exclusion

This specification uses the term included to mean traffic forwarded through the L2VPN forwarder, and excluded to mean all other non-L2VPN traffic.

³⁷ Sentence modified per L2VPN-N-12.1065-11 on 9/19/12 by CP.

³⁸ Sentence modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.

Subscriber Transparent LAN Service requires that all CPE traffic be included in L2VPN forwarding while traffic from embedded CMs and any other embedded hosts co-located with the CM be excluded from L2VPN forwarding. The L2VPN feature can be configured, however, so that eCM and eSAFE traffic can be forwarded on separate L2VPNs if desired.

6.6.1 CM and eSAFE Host Forwarding Model

Figure 6–1 depicts the overall CMTS and CM L2VPN forwarding model for CMs and eSAFE Hosts.

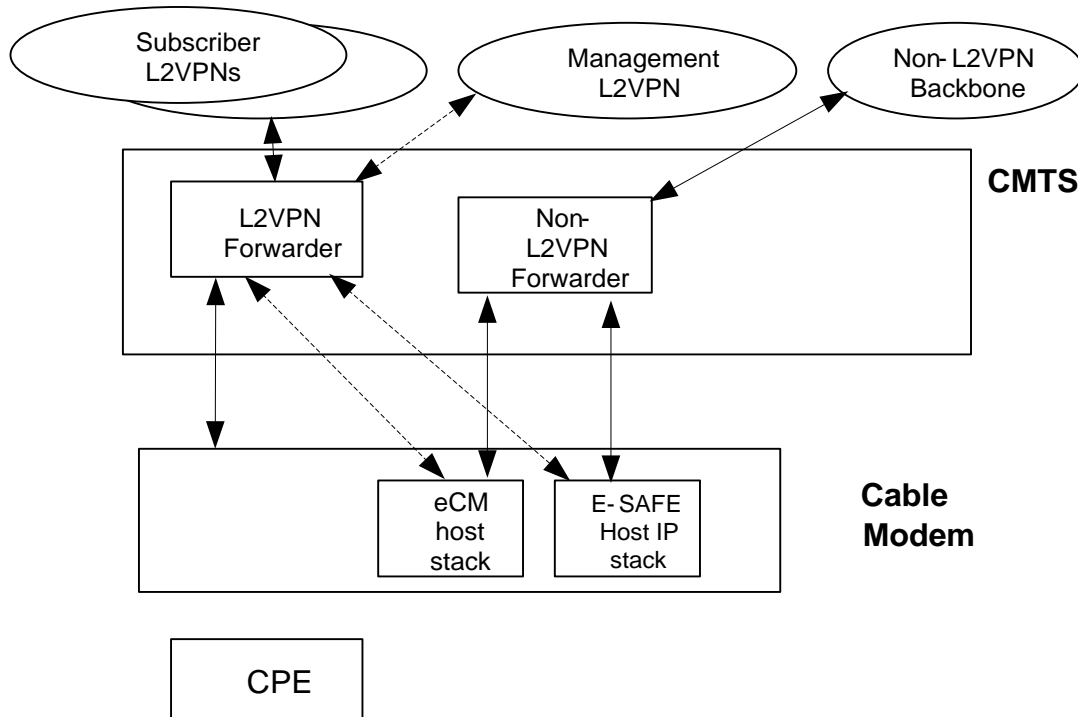


Figure 6–1 - CM, eMTA, and CPE Forwarding

For TLS service, traffic to and from the eCM's and eSAFE's host IP stack is excluded from L2VPN forwarding, and is forwarded with the CMTS's normal non-L2VPN forwarder, as depicted with a solid line in Figure 6–1.

With the Management L2VPN feature, the eCM's post-registration traffic is classified to an L2VPN Forwarding service flow and included for L2VPN forwarding (dotted line). Any pre-registration eCM traffic cannot use the L2VPN feature because encryption in an L2VPN SAID is not possible prior to registration.

The management traffic to and from an eSAFE host can use the same Management L2VPN as the eCM, or its own separate Management L2VPN, as desired.

6.6.2 Cable Modem MAC Bridge Interface Masks

In accordance with the MAC bridging model described in [eDOCSIS], a CM implementing this specification is considered to implement a set of MAC Bridge Interfaces, as summarized below.

Table 6–2 - Cable Modem MAC Bridge Interfaces ³⁹

ifIndex	Description
(0)	(eCM: self-Host interface)
1	Primary CPE Interface, also eRouter: DOCSIS Embedded Router
2	RF Interface
16	eMTA: PacketCable embedded Multimedia Terminal Adapter host interface
17	eSTB-IP: OpenCable embedded Set Top Box IP Host interface
18	eSTB-DSG: OpenCable embedded Set Top Box DOCSIS Set-top Gateway interface

This specification introduces the convention that ifIndex 0 is considered by convention to apply to the internal host interface to the CM's management IP stack, or its self interface.

L2VPN layer 2 forwarding in a CM is considered to occur only on an explicit list of the above MAC Bridge interfaces. For example, Transparent LAN Service involves bridging only from the RF MAC Interface (ifIndex 2) to the Primary CPE Interface (ifIndex 1). TLS is not permitted to access the eCM's self interface or any other eSAFE host interface.

For each L2VPN that is forwarded within a CM, a CM Interface Mask (CMIM) parameter is configured with the set of MAC Bridge interfaces that are permitted to forward packets to and from that L2VPN. Each MAC Bridge interface is assigned a bit position in the CMIM mask corresponding to its ifIndex value. The eCM's host interface is assigned CMIM bit position 0.

The CMIM parameter for an L2VPN is encoded in the same per-CM L2VPN encoding that defines the L2VPN VPNID. If the CMIM subtype is omitted from a forwarding L2VPN encoding, its default value is the one appropriate for TLS service, (i.e., with only the RF Interface (ifIndex 2) and Primary CPE Interface (ifIndex 1) bits set). The CMIM parameter is encoded in the same manner as the Basic Encoding Rules of an SNMP BITS object type. In a CMIM Subtype TLV, the bit mask is encoded as a variable length octet string where bit position 0 is the most significant bit of the first octet; position 1 is the next most-significant bit; position 7 is the least significant bit of the first octet. Bit position 8 is the most significant bit of the second octet. As an example, the default CMIM value, with bit positions 1 and 2 set, can be encoded as a single octet with the value 0x60.⁴⁰

6.6.3 Embedded Host Exclusion

When a CMIM mask has a value of zero in a MAC bridge interface position, all traffic from that interface is configured to be excluded from L2VPN forwarding. In particular, the eCM self host interface (CMIM interface bit position 0) is excluded from Transparent LAN Service L2VPNs. When the eCM self host interface is excluded from a CMIM subtype for an L2VPN forwarding upstream SF, the CMTS MUST exclude from upstream L2VPN forwarding all traffic that contains a source MAC that matches the CM host's MAC address.

Because non-compliant CMs are unable to classify L2VPN from non-L2VPN upstream traffic, a CMTS MUST support both L2VPN and non-L2VPN forwarding of upstream traffic from a Forwarding L2VPN service flow of a non-compliant CM based on checking the source MAC address against a CM Interface Mask configured for the SF. The CMTS MUST direct packets from included host types to the L2VPN forwarder. The CMTS MUST NOT deliver traffic from excluded host types to the L2VPN Forwarder.

A CMTS MUST support exclusion of the CM MAC address and at least one other eSAFE MAC address on all L2VPN forwarding service flows from both non-compliant CMs and compliant CMs. Non-compliant CMs may have at most, one eSAFE MAC address checked in this manner. This specification does not support L2VPN forwarding from a non-compliant CM with more than one eSAFE host. This specification supports L2VPN

³⁹ Table revised per L2VPN-N-12.1065-11 on 9/19/12 by CP. Revised per L2VPN-N-14.1135-5 on 3/27/14 by JB.

⁴⁰ Paragraph modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.

forwarding from compliant CMs with more than one eSAFE host only by configuring Upstream Packet Classification encodings that explicitly classify the non-L2VPN forwarding eSAFE traffic to non-L2VPN forwarding upstream service flows.

6.6.4 CMTS embedded host MAC Address Learning

A CMTS MUST learn the MAC address of an embedded CM from the Source MAC address of the CM's initial ranging request message and include it in the docsDevCmCmtsStatusTable.

The CMTS uses two techniques to learn the MAC address of eSAFE hosts:

- For L2VPN-compliant CMs, the CMTS MUST learn the eSAFE MAC addresses from the eSAFE Host Capability encodings Section B.1.2 when the CM registers;⁴¹
- For non-L2VPN-compliant CMs, the CMTS MUST snoop upstream DHCP packets to determine the eSAFE MAC addresses.

6.6.4.1 Enable eSAFE DHCP Snooping Subtype⁴²

The CMTS MUST enable DHCP snooping to determine eSAFE MAC addresses from a non-compliant CM when an Enable eSAFE DHCP Snooping subtype is present in any per-SF L2VPN Encoding (Section B.3.3). The value of the encoding is a bit mask that enables particular eSAFE hosts to be snooped. The CMTS MUST NOT enable DHCP snooping when the Enable eSAFE DHCP Snooping subtype is absent from all per-SF L2VPN encodings or does not have a '1' bit for the particular eSAFE host type when it is present. This is to prevent spoofing of eSAFE by unauthorized non-embedded CPEs.

When eSAFE DHCP snooping is enabled, the CMTS MUST support detection of the eSAFE host type of a MAC address, from the initial substring of option 60 of the broadcast DHCP DISCOVER packet, from the eSAFE host that is relayed by the CMTS, when option 60 is present. When eSAFE DHCP snooping is enabled, the CMTS MUST support detection of the eSAFE host type of a MAC address from option 43, subtype 2 of a DHCP DISCOVER packet from the eSAFE host that is relayed by the CMTS, when option 43, subtype 2 is present. Table 6–3 provides the values of these DHCP options for each currently defined eSAFE host type.

Table 6–3 - eSAFE DHCP Snooping Substrings⁴³

ESAFE Host Type	DHCP Option 60 substring	DHCP Option 43, sub-option 2 substring
Embedded Multimedia Terminal Adapter [PKT-PROV]	pktc	EMTA
eRouter [eRouter]	eRouter1.0	EROUTER

The CMTS MUST learn the eSAFE's MAC address from the client hardware identifier field of the snooped DHCP-DISCOVER packet.

Once the CMTS learns the MAC addresses of eSAFE hosts, the CMTS thereafter excludes from L2VPN forwarding, all upstream traffic from that DOCSIS host MAC address.

6.6.5 Interface-based Classification

The RF MAC Domain implements the DOCSIS Upstream Packet Classifiers that classify an L2PDU bridged to the MAC Domain interface into an upstream Service Flow.

In an Upstream Packet Classification encoding, the CMIM subtype represents a rule that matches the ingress bridge port of the L2PDU. This allows classifiers to classify CPE, eCM, and eSAFE traffic generically, by host type, instead of requiring static classifiers to be based on the actual MAC or assigned IP address of the host.

The Host Classification capability is most useful when implementing Management L2VPNs for isolating CM and eSAFE management traffic from payload traffic on a layer 2 backbone. It allows the CM's classifiers to classify

⁴¹ Paragraph modified per L2VPN-N-07.0461-2 on 7/5/07 by KN.

⁴² Paragraph modified per L2VPN-N-07.0461-2 on 7/5/07 by KN.

⁴³ Table revised per L2VPN-N-12.1065-11 on 9/19/12 by CP. Revised per L2VPN-N-14.1135-5 on 3/27/14 by JB.

upstream eCM and/or eSAFE traffic to a separate L2VPN forwarding upstream service flow for the Management L2VPN.

6.7 L2VPN Quality Of Service ⁴⁴

An operator may offer L2VPN services with Service Level Agreements (SLAs) inclusive of Quality of Service (QoS). This may require QoS mechanisms not only on DOCSIS access, but also on the backbone network. Section 6.7.1 describes the means for DOCSIS QoS, and Section 6.7.2 describes the means for Backbone Network QoS.

This specification describes QoS from the perspective of the DOCSIS network. It corresponds with the Metro Ethernet Forum (MEF) and IEEE 802 terms of ingress and egress as follows:

- Upstream corresponds with the MEF/IEEE 802 term of ingress from the perspective of the UNI.
- Downstream corresponds with the MEF/IEEE 802 term of egress from the perspective of the UNI.

6.7.1 DOCSIS QoS

6.7.1.1 Service Flow Separation

An important aspect of the L2VPN service offering by an operator is isolation not only of traffic forwarding, but also of Quality of Service. It should not be possible for one L2VPN traffic flow (or even a non-L2VPN traffic flow) with excessive traffic to significantly affect the QoS received by any other L2VPN's flows. Accordingly, this specification requires that the downstream traffic for each L2VPN be isolated from each other and from non-L2VPN traffic by placing the traffic in a separate Service Flow (SF). In the case of Point-to-Point mode forwarding, this happens automatically because each CM already has a primary downstream service flow. In the case of Multipoint mode forwarding, this specification requires each L2VPN to have a separate service flow for its downstream flooded group MAC and unknown individual MAC destined traffic.

6.7.1.2 Classification and Scheduling

DOCSIS QoS mechanism defined in [DOCSIS MAC] can be utilized for classification and scheduling of L2VPN traffic. This specification allows leveraging existing Layer 2 criteria for upstream and downstream service flow classification, and defines an additional (new) criterion for downstream service flow classification (see Section 6.7.2.2).

Suffice to say, only the QoS Parameter Set of the service flow, to which the L2VPN traffic is classified, defines the forwarding priority of L2VPN traffic on a DOCSIS RF interface in both upstream and downstream directions.

6.7.2 Backbone Network QoS

If L2VPN traffic is forwarded and received over the backbone network (Ethernet or IP or MPLS), then prioritized forwarding (over best effort, say) may be desired to ensure appropriate QoS for L2VPN traffic in the backbone network. The prioritized forwarding typically relies on the priority encoded in the L2VPN packet itself, where the priority may be indicated in a variety of ways, depending on the chosen NSI Encapsulation:

- In the user-priority bits of an outer IEEE 802.1 tag (see Section 6.7.2.1.1);
- In the Traffic Class bits (formerly known as EXP experimental bits) of an MPLS label (see Section 6.7.2.1.2);
- In the IP Precedence bits of an L2TPv3 pseudowire encapsulation (see Section 6.7.2.1.3).

The CMTS MUST transmit an L2VPN packet upstream on an NSI port with the priority encoded as appropriate for its NSI encapsulation. In other words, if CMTS forwards the encapsulated L2PDU on the NSI, then it MUST appropriately set the priority of the encapsulated L2PDU to help with upstream QoS, if any, on NSI and beyond. Similarly, if CMTS receives the encapsulated L2PDU on the NSI, then it may inspect the priority of the encapsulated L2PDU to help with downstream QoS on the DOCSIS RFI.

This specification enables supporting not only Short Pipe model, but also Pipe model, see [RFC 3270].

⁴⁴ Section and subsections modified per L2VPN-N-12.1065-11 on 9/19/12 by CP.

6.7.2.1 User Priority

6.7.2.1.1 IEEE 802.1 User Priority

The IEEE 802.1 layer 2 bridging model uses the concept of a user priority with 8 possible values to indicate the QoS to be provided when forwarding an L2PDU [802.1Q]. This field is used to provide differentiated QoS to different traffic flows in the Ethernet Network.

6.7.2.1.2 MPLS Traffic Class

The MPLS uses the concept of a Traffic Class (TC), formerly known as EXP, with 8 possible values to indicate the QoS to be provided when forwarding an MPLS packet carrying payload such as L2PDU. This TC field is used to provide differentiated QoS to different traffic flows in the MPLS network.

6.7.2.1.3 IP Precedence

The IP uses the concept of a DSCP bits with 8 or more possible values to indicate the QoS to be provided when forwarding an L2TPv3 over IP packet carrying payload such as L2PDU. This field is used to provide differentiated QoS to different traffic flows in the IP network.

6.7.2.2 Downstream User Priority Range Classification

This specification defines a Downstream User Priority Range subtype of an L2VPN Encoding that appears as a new rule matching criterion in a downstream Service Flow Packet Classifier encoding. The Downstream User Priority Range subtype is described in Section B.3.9 and is intended to support Pipe Model QoS [RFC 3270]. When a Downstream User Priority Range subtype is present in a Downstream Service Flow Packet Classifier Encoding, the CMTS MUST match the classifier only to L2VPN forwarded L2PDUs with a user priority (as defined in Section 6.7.2.1) within the indicated range. In the case of Multipoint L2VPN forwarding, the CMTS MUST compare the user priority set by the Upstream User Priority subtype on traffic received by the CMTS, with the Downstream user priority range subtype in the Downstream classifier encoding when hairpinning the traffic downstream. This subtype enables matching on the priority of the encapsulated L2PDU as it is received on the NSI by CMTS, to help with the downstream QoS on DOCSIS RF interface. The CMTS MUST forward unmatched packets on the primary downstream service flow of the CM.

6.7.2.3 Downstream User Priority

A CMTS MUST accept the priority bits of a service-delimiting 802.1Q tag on the NSI port as the L2VPN user priority attribute of the packet. A CMTS MUST accept the TC field of the outermost MPLS label received on the NSI port as the L2VPN user priority attribute of the packet. A CMTS MUST accept the DSCP bits of an L2TPv3 encapsulated IP packet received on the NSI port as the L2VPN user priority attribute of the packet. The CMTS MUST maintain the user priority of the packet within the L2VPN forwarder (possibly regenerating it via vendor-specific configuration), and use this value for matching against the Downstream User Priority Range subtype of Downstream Classifier L2VPN Encodings.

6.7.2.4 Upstream User Priority⁴⁵

This specification defines an Upstream User Priority subtype of an L2VPN Encoding that appears in an upstream Service Flow. The Upstream User Priority subtype is described in Section B.3.8.

This specification calls for the CMTS to set a configured user priority on each NSI encapsulated L2VPN packet for an upstream L2VPN forwarding service flow based on a configured Upstream User Priority subtype of the per-SF L2VPN Encoding that defined the SF. The Upstream User Priority subtype essentially defines the priority of the L2PDUs forwarded by CMTS on NSI across the cable operator backbone.

When an L2PDU is forwarded with an IEEE 802.1Q Tag on a trunk Ethernet NSI interface of CMTS, the priority of the packet is encoded in the upper 3 bits of an 802.1Q Tag Control value. Appendix II provides details of the IEEE 802.1Q encapsulation.

⁴⁵ Updated per L2VPN-N-14.1186-1 on 1/15/15 by PO.

When an L2PDU is forwarded over an MPLS Pseudowire on an IP/MPLS enabled NSI interface of CMTS, the priority of the packet is encoded in MPLS TC field (3 bits). The CMTS encodes the priority into both the outer and inner MPLS labels to support both Short Pipe and Pipe mode operation.

When an L2PDU is forwarded over an L2TPv3 Pseudowire on an IP enabled NSI interface of CMTS, the priority of the packet is encoded in first 3 bits of IP DSCP (i.e., IP Precedence bits).

A CMTS MAY implement vendor-specific configuration to set the upstream user priority of NSI encapsulated L2VPN packets. The particular bridging model implemented by the CMTS (e.g., [802.1Q], or MPLS) MAY provide for the regeneration of a different user priority for the packet in the CMTS from that which is sent by the CPE.

The CMTS MUST NOT use any 802.1 priority bits applied by CPE to determine the user priority of an NSI encapsulated packet. If necessary, the CM can be configured to classify the upstream priority-only CPE-tagged packet to a service flow that is configured with an explicit Upstream User Priority subtype. If the CMTS implements an IEEE 802.1Q bridge forwarder, the CMTS MAY map the inner customer user priority tag (802.1P) to the outer service user priority tag (802.1P).

NOTE: The user priority parameter of an L2VPN packet defines only the priority of the packet's forwarding across the cable operator backbone; it does not affect the forwarding of the packet upstream or downstream on the DOCSIS RF interface. Only the QOS Parameter Set of the service flow to which the packet is classified defines the priority for forwarding the packet on a DOCSIS RF interface.

The CMTS MUST set the user priority as explicitly configured by the Upstream User Priority subtype. If the Upstream User Priority subtype is not set, then the CMTS MUST set the user priority to 0. The L2VPN forwarder in the CMTS MUST set the appropriate user priority (possibly regenerated) of the IEEE 802.1Q tag that is imposed on the L2PDU prior to forwarding it out of an NSI port. The L2VPN forwarder in the CMTS MUST set the appropriate user priority (possibly regenerated) of the MPLS TC field that is imposed on the L2PDU prior to forwarding it out of an NSI port. The L2VPN forwarder in the CMTS MUST set the appropriate user priority (possibly regenerated) of the DSCP bits of the L2TPv3 encapsulated IP packet that is imposed on the L2PDU prior to forwarding it out of an NSI port.

6.8 Stacked 802.1Q Tags or Tag-in-Tag operation ⁴⁶

The selection of the particular L2VPN for upstream bridged traffic is always indicated by the VPNID subtype of the L2VPN Encoding. This specification does not address the interpretation of 802.1Q tags applied by CPE and received by the Ethernet port of a CM for forwarding upstream. Such tags are considered non-service delimiting, and are always ignored for purposes of L2VPN selection by the DOCSIS CMTS. These non-service delimiting tags are forwarded as part of the CPE payload. DOCSIS CMTSs and CMs conforming to this specification support forwarding of maximum length Ethernet packets with a single subscriber non-service delimiting 802.1Q tag. This means that a pre-3.1 DOCSIS CM supporting this specification is able to forward 1522 byte packets between its RF and CPE interfaces and a CMTS is able to forward 1526 byte packets on its 802.1Q tagging Ethernet NSI port. A DOCSIS 3.1 CM is able to forward 2000 byte packets between its RF and CPE interfaces and a CMTS is able to forward 2004 byte packets on its 802.1Q tagging Ethernet NSI port.

When a packet with an inner CPE-supplied non-service delimiting tag is forwarded onto an NSI port that is also using IEEE 802.1Q Encapsulation, the CMTS adds an outer service-delimiting tag with the configured NSI Encapsulation VLAN ID. This is called stacked or tag-in-tag 802.1Q tagging. The IEEE 802.1p/Q Packet Classification Encoding criteria in [DOCSIS MAC] apply to only the non-service delimiting CPE-supplied 802.1Q tag as the packet appears on the RF interface, not the outer service-delimiting tag value as the packet appears on an NSI port. For example, the CM may classify upstream CPE packets to a particular service flow, based on the priority bits of the CPE-applied tag.

In order to prevent CPE abuse of the backbone network, the CMTS MUST NOT interpret a priority-only tag applied by the CPE as defining the upstream user priority of an L2VPN packet. The user priority of an upstream packet is defined only by the configured Upstream User Priority subtype of the per-SF Forwarding L2VPN Encoding that applies to the packet. A CPE-applied priority-only tag is treated as a non-service delimiting tag and is stacked as an inner tag when forwarded by the CMTS. If a CPE-applied priority-tag is desired to select the upstream user priority,

⁴⁶ Section modified per L2VPN-N-07.0471-2 on 7/5/07 by KN. Revised per L2VPN-N-14.1135-5 on 3/27/14 by JB.

the CM should be configured to classify the packet to a service flow with an explicit Upstream User Priority subtype. This allows the cable operator to control the priority of forwarded L2VPN packets on the backbone.⁴⁷

In the downstream direction, the IEEE 802.1p/Q Packet Classification Encoding criteria of [DOCSIS MAC], apply only to any inner, non-service delimiting tag in the packet as it appears on the RF interface.

NOTE: CMTSs are not required to implement these layer 2 criteria in the downstream direction. What is usually desired, however, is to classify downstream traffic according to the priority or VLAN ID of the outer service delimiting tag as the packet appeared on the NSI interface. This specification defines Downstream Classifier L2VPN Encodings to permit classification based on the packet's VPNID and user priority as signaled in its NSI encapsulation.

6.9 Spanning Tree and Loop Detection⁴⁸

[DOCSIS MAC] describes the DOCSIS Spanning Tree Protocol (DSTP). Because of limited implementation, this protocol cannot be relied upon to avoid L2VPN bridging loops. An operator is expected to configure subscriber L2VPN networks in a loop-free manner, or to rely upon subscriber equipment itself implementing the IEEE Spanning Tree Protocol to break any bridging loop on a subscriber's L2VPN. This section describes CMTS requirements to prevent L2VPN denial of service when a subscriber accidentally or intentionally configures a bridging loop.

The CMTS MUST transparently forward the IEEE Spanning Tree Protocol (STP) on the subscriber's layer2 VPN.

A CMTS MAY implement the DOCSIS Spanning Tree protocol and transmit DSTP BPDUs on all NSI and RF interfaces configured for L2VPN operation. The CMTS MUST transmit DOCSIS Spanning Tree Protocol packets as untagged on an IEEE 802.1Q NSI interface and encrypted in a SAID provided to the L2VPN CMs on a CMTS MAC domain RF interface. A CMTS MAY implement a DOCSIS Spanning Tree (DST) SAID specifically for DST forwarding to the CPE ports of all L2VPN CMs.

A CMTS MUST prevent a bridging loop for one L2VPN from denying all forwarding or flooding of traffic on any other non-looped L2VPN. A CMTS MAY require configuration of both downstream and upstream service flow maximum forwarding rates to meet this requirement, as long as such limits are no less than 10% of link capacity.

6.10 High Availability⁴⁹

This section describes how a CMTS provides high availability of L2VPN service in case of a link failure on the NSI.

6.10.1 802 Encapsulation - Active/Standby Layer 2 Trunk Ports

The CMTS can provide high availability forwarding for layer 2 encapsulated L2VPN packets by permitting a single "active" layer 2 NSI trunk port and one or more "standby" layer 2 trunk ports on a per-L2VPN basis. The CMTS MUST implement configuration to enable or disable an NSI port for "L2 Trunk Port" operation, with a default of 'disabled'. The CMTS MUST implement configuration of an 'L2 Trunk Priority' attribute for an NSI port in the range 0..255 with a default value of 128, where higher priority values are more preferred. The CMTS MUST select for "active" forwarding of IEEE 802 layer 2 encapsulated L2VPN packets an operational NSI port enabled for L2 trunk port operation with a configured trunk priority no lower than any other operational NSI port enabled as an L2 trunk port.

The CMTS MAY support concurrent forwarding of layer 2 encapsulated L2VPN traffic across multiple NSI ports with the same trunk priority. In this case, the CMTS MUST enforce that a given VLAN ID is forwarded upstream and accepted for forwarding downstream only from a single such port. When the CMTS does not support concurrent layer 2 encapsulated forwarding through same-priority trunks, the CMTS MUST reject configuration of multiple NSI trunk ports with the same layer 2 trunk priority. Note that configuration of the layer 2 trunk port operation or priority attribute does not affect forwarding or acceptance of L2VPN traffic with encapsulations other than IEEE 802 layer 2.

⁴⁷ Paragraph modified per L2VPN-N-12.1065-11 on 9/19/12 by CP.

⁴⁸ Section modified per L2VPN-N-07.0471-2 on 7/5/07 by KN.

⁴⁹ Section added per L2VPN-N-12.1065-11 on 9/19/12 by CP.

6.10.2 802 Encapsulation - Link Aggregation

The CMTS MUST support configuration and operation of L2VPN forwarding with any IEEE 802 layer 2 encapsulation across a link aggregation group of NSI ports operating according to [802.1AX].

The CMTS MAY support Link Aggregation Control Protocol. The CMTS MUST support configuration to operate without Link Aggregation Control Protocol.

6.10.3 MPLS Encapsulation High Availability

The CMTS MUST support configuration and operation of L2VPN MPLS forwarding on more than one NSI ports. The CMTS MUST support standard MPLS control plane protocols (e.g., LDP) and mechanisms (e.g., LDP re-convergence) for re-establishing the L2VPN forwarding on the alternative NSI port when an NSI port fails.

6.11 MPLS Encapsulation - PW Redundancy ⁵⁰

The CMTS MUST support configuration and operation of L2VPN forwarding using a backup MPLS Peer and backup pseudowire ID. The backup MPLS Pseudowire is active when the primary has failed.

- Backup MPLS Peer: the backup MPLS peer is used to define a second Pseudowire destination IP for use in pseudowire redundancy.
- Backup Pseudowire ID: this is the backup MPLS Pseudowire ID also needed to define the second pseudowire endpoint for use in pseudowire redundancy.

⁵⁰ Section added per L2VPN-N-12.1065-11 on 9/19/12 by CP.

7 CABLE MODEM REQUIREMENTS ⁵¹

A CM MUST accept one or more L2VPN SA-Descriptor Subtypes added by a CMTS to any forwarding L2VPN Encoding in a REG-RSP, REG-RSP-MP, or DSx-RSP message to the CM. The CM associates the SAIDs of the SA-Descriptors in the L2VPN Encoding with the single L2VPN identified in the L2VPN Encoding. The CM MUST be capable of associating any number of its available SAIDs to an L2VPN. The CM MUST be capable of associating more than one SAID to a single L2VPN. A CM receiving an L2VPN Encoding with an L2VPN SA-Descriptor Subtype for a SAID not previously established on that CM MUST initiate a BPKM TEK transaction to establish the new L2VPN SAID [BPI+]. A CM receiving an L2VPN SA-Descriptor Subtype in a REG-RSP or REG-RSP-MP MUST wait for BPI Authorization to complete before initiating the BPKM TEK. The CM determines that a downstream packet is to be forwarded on an L2VPN when the packet is encrypted with an L2VPN SAID. A CM MUST replace the set of L2VPN SAIDs of an L2VPN when receiving a top-level L2VPN Encoding in a MAC Management message that identifies that L2VPN. The CM MUST discontinue downstream decryption of an L2VPN SAID when it receives in a dynamic service flow message a top-level L2VPN Encoding for an L2VPN ID that omits the SA-Descriptor subtype with that SAID.

A CM MUST promiscuously forward all downstream group MAC (GMAC) destined traffic that is encrypted in an L2VPN SAID signaled to the CM, regardless of the GMAC destination of the packet.

The CM MUST NOT apply the multicast filtering or forwarding rules of [DOCSIS MAC] to downstream GMAC traffic encrypted in an L2VPN SAID. A CM MUST continue to implement downstream DOCSIS group MAC forwarding rules for all unencrypted packets and packets encrypted in a non-L2VPN SAID.

A CM MUST NOT implement the IGMP Multicast Forwarding rules of [DOCSIS MAC] for any upstream packets (e.g., IGMP Membership Reports) classified to a forwarding L2VPN service flow. The CM MUST continue to implement DOCSIS IGMP Multicast Forwarding rules for upstream IGMP Membership Reports not classified to a forwarding L2VPN service flow.

A CM MUST restrict bridge forwarding of downstream packets encrypted in an L2VPN SAID to only the bridge interfaces indicated with a '1' bit in the CM Interface Mask (CMIM) configured for that L2VPN.⁵² For example, the CM does not deliver downstream L2VPN traffic to the eCM or any eSAFE internal host if the CMIM omits that host interface (i.e., contains a '0' bit for that interface), even if the packet is addressed to the individual destination MAC address for the host. Likewise, the CM does not deliver L2VPN-labelled group MAC (GMAC) destined traffic to internal hosts when that L2VPN's CMIM omits the internal host interface.

A CM MUST forward downstream packets encrypted in an L2VPN SAID with a known destination MAC address (DMAC) to the L2VPN interface with which the destination MAC address is associated. The DMAC may be associated via a configured CMIM or learned via traffic received from L2VPN CPEs. The CM MUST forward downstream packets encrypted in an L2VPN SAID with an unknown DMAC to all L2VPN interfaces other than the interface on which it was received. When forwarding of frames received from any CPE interface, the CM forwards the frames addressed to unknown destination MAC addresses only to the RF Interface as per Section on 'CM Operational Forwarding Behavior' in [MULPIv3.0].⁵³

A CM MUST support a classification rule criterion signaled with a CM Interface Mask (CMIM) in an L2VPN Encoding of an Upstream Classifier Packet Encoding, whether or not that Encoding classifies to a Forwarding L2VPN service flow. The CM MUST consider the criterion to be matched when the source MAC address of an upstream packet is for a host type with a '1' bit in the CM Interface Mask. The CM MUST consider the criterion to be unmatched and forward or drop a packet accordingly, when the source MAC address is for a host type with a '0' bit in the CM Interface Mask.

A CM MUST support the Downstream Unencrypted Traffic (DUT) Filtering feature as described in Section B.2, and advertise this in a DUT Filtering Capability Encoding, Section B.1.3. When DUT Filtering is enabled, a CM MUST restrict bridge forwarding of downstream unencrypted traffic to only the interfaces indicated in the DUT CM Interface Mask (DUT CMIM) implied or configured by the DUT Filtering Encoding.

⁵¹ This section modified per L2VPN-N-07.0471-2 on 7/5/07 and L2VPN-N-08.0625-2 on 4/16/08 by KN. Revised per L2VPN-N-14.1135-5 on 3/27/14 by JB.

⁵² Sentence modified per L2VPN-N-12.1065-11 on 9/19/12 by CP.

⁵³ Paragraph added per L2VPN-N-12.1065-11 on 9/19/12 by CP.

Because CMIM bit position 1 (corresponding to CM bridge ifIndex 1) represents the *set of all* CPE interfaces in L2VPN Forwarding, DUT Filtering, and Upstream Classifier Encodings, a compliant CM that implements more than one CPE interface MAY assign a CMIM bit position in the range of 5..15 to represent its single primary CPE interface. This is so that CMIM values (and other DOCSIS interface-specific filters) can represent the primary CPE interface by itself, independent of the set of all other CPE interfaces. The CM MUST continue to report only ifIndex 1 as its primary CPE interface.

A CM MUST forward upstream and downstream packets as large as the maximum supported packet PDU length as defined in [DOCSIS MAC].

A CM SHOULD support acquisition of at least 128 CPE MAC addresses, indicated in the Max CPE encoding in the config file as defined in [DOCSIS MAC].

A CM MUST support L2VPN MAC Aging Mode, described in Section B.7. When L2VPN MAC Aging Mode is enabled, the CM bridging table is full, and traffic destined for the upstream RF interface arrives from a previously unknown MAC address, the CM will overwrite an entry in the bridging table that corresponds to the device that has not transmitted a packet destined for the upstream RF interface for the longest amount of time with the new MAC address.

A CM MUST advertise the L2VPN Capability subtype of the Modem Capabilities Encoding (Section B.1.1) of its Registration Request.⁵⁴

A CM with embedded eSAFE hosts MUST advertise them to the CMTS in a Registration Request message with an eSAFE Host Capability Encoding (Section B.1.2) for each eSAFE host.⁵⁵

A CM MUST silently ignore an L2VPN Encoding in any TLV context not mentioned in this specification. A CM MUST silently ignore any unrecognized L2VPN Encoding subtype and process all recognized L2VPN Encodings normally.

A CM MUST support classification of traffic based on the IEEE 802.1P/Q Packet classification encodings [MULPIv3.0].⁵⁶

⁵⁴ Paragraph revised per L2VPN-N-07.0461-2 on 7/5/07 by KN.

⁵⁵ Paragraph revised per L2VPN-N-07.0461-2 on 7/5/07 by KN.

⁵⁶ Added per L2VPN-N-12.1065-11 on 9/19/12 by CP.

8 SERVICE OPERATIONS, ADMINISTRATION, AND MAINTENANCE (OAM)⁵⁷

8.1 Introduction

Service OAM is a general term for the management capabilities necessary to install, monitor, and troubleshoot an Ethernet service and allow service providers to offer predictable levels of service. Of interest to MSOs, Service OAM consists of a series of techniques for fault management and performance management. Fault management capabilities include the ability to detect, verify, and isolate faults in a network. These techniques allow operators to detect and isolate both hard failures such as circuit outages and soft failures such as network element misconfiguration or software failure. Performance management capabilities include the ability to measure frame delay, frame delay variation, frame loss, and service availability. These techniques allow operators to efficiently meet performance targets across a network while minimizing operational expenses.

The Service OAM requirements in this specification define a "lightweight" subset of Fault Management and Performance Management functions described in the following subsections. Neither the CM nor the CMTS are required to provide full support for Service OAM requirements as described by [802.1Q], [Y.1731], [MEF 30.1] and [MEF 35].

The DOCSIS network transmits Service OAM traffic as part of an L2VPN. Therefore, such traffic receives the quality of service configured for the L2VPN.

Table 8–1 identifies the SOAM FM and PM requirements that are specified in this specification. The table describes a high level role of CM and CMTS with each of the requirements and also provides the original references to [802.1Q] and [Y.1731].

Table 8–1 - High Level Operations and References.

Requirements	CM/CMTS Role	References
Continuity Check Message	CM (acting as MEP) to send and receive CCM. CMTS (acting as MEP) to send and receive CCM.	Clause 21.6 [802.1Q] Clause 7.1, 9.2 [Y.1731]
Loopback	CM (acting as MEP) to respond to LBM with LBR. CMTS (acting as MIP) to respond to LBM with LBR. CMTS (acting as MEP) to initiate and send LBM toward CM.	Clause 21.7 [802.1Q] Clause 7.2, 9.3, 9.4 [Y.1731]
Linktrace	CM (acting as MEP) to respond to LTM with LTR. CMTS (acting as MIP) to respond to LTM with LTR. CMTS (acting as MIP) to forward the LTM toward the remote MEP. CMTS (acting as MEP) to initiate and send LTM toward CM.	Clause 21.9 [802.1Q] Clause 7.3, 9.5, 9.6 [Y.1731]
Frame Delay	CM (acting as MEP) to respond to DMM with DMR. CMTS (acting as MEP) to initiate and send DMM toward CM.	Clause 8.2, 9.15, 9.16 [Y.1731]
Synthetic Frame Loss	CM (acting as MEP) to respond to SLM with SLR. CMTS (acting as MEP) to imitate and send SLM toward CM	Clause 8.4, 9.22, 9.23 [Y.1731]

⁵⁷ New section added per L2VPN-N-12.1065-11 on 9/19/12 by CP. Revised per L2VPN-N-14.1135-5 on 3/27/14 by JB.

8.2 Service OAM Configuration⁵⁸

Table 8–2 shows Service OAM terminology mapping between [802.1Q] and [Y.1731].

Table 8–2 - Service OAM Terminology

IEEE 802.1Q Terminology	ITU-T Y.1731 Terminology
Maintenance Association (MA)	Maintenance Entity Group (MEG)
Maintenance Association Identifier (MAID)	Maintenance Entity Group Identifier (MEG ID)
Maintenance Domain (MD)	(no such construct)
Maintenance Domain Level (MD Level)	Maintenance Entity Group Level (MEG Level)

In this document, we will use terminology from ITU-T except when discussing Maintenance Association Identifier and Maintenance Domain.

OAM frames belonging to an administrative domain originate and terminate in Maintenance Entity Group (MEG), End Points (MEPs) [Y.1731], present at the boundary of that administrative domain. The MEPs associated with administrative domain operate at an assigned level. The CM operates as a MEP. Cable Modem interfaces will be considered terminating interfaces for the subscriber Ethernet service. Thus, they will not have to serve as interfaces to another provider. The MEP is configured on a per L2VPN basis.

The MEG Intermediate Points (MIPs) [Y.1731] reside in the middle of a network. They generally do not originate OAM traffic but only respond to OAM PDUs. CMTSs operate as MIPs, although they can also operate as MEPs.

The Service OAM traffic gets the same QoS as the other L2VPN traffic, including the Upstream User Priority set for the L2VPN service flow.

By default, all SOAM Fault Management and Performance Management functions described in Sections 8.3 and 8.4 are disabled unless they are explicitly enabled per TLV configuration Annex B.

8.2.1 CM

The CM is configured with a MAID, MEP ID, List of Peer MEP ID(s), and MEG level.

The MAID/MEG ID is an identifier for the MEG and it's required to be unique across the operator's network. The MAID is formed with the Maintenance Domain (MD) Name and Short Maintenance Association (MA) Name. For MD, this document recommends using Maintenance Domain not present (value 1 in table 21-18 of [802.1Q]) and Short MA Name Format value of 2, which is character string. Other options for Short MA Name Format exist and can be found on table 21-20 of [802.1Q].

The MAID, MEP ID and MEG levels are configured in the CCM frame format. The detailed structure can be found in Clause 21 [802.1Q].

The MEP ID is an integer in the range of 1..8191 and it is unique ID among all the MEPs in a MEG.

The CM MUST support configuration as a MEP using the configuration TLVs defined in Annex B.

The CM might be configured with a MEG level (0..7). See [MEF 30.1] section 7.1 for default values for each MEG level. The CM MUST support UP MEPs at two MEG levels, Operator ME (default level 2) and EVC ME (default level 4). The CM MAY support other MEG levels (0..7) as defined in [MEF 30.1]

A CM MAY support a MEP instance for each configured L2VPN. A CM could be configured with multiple L2VPNs where some L2VPNs could have a MEP configured, while others do not.

Figure 8-1 describes the MEG level that will be supported by the CM and the CMTS. The full diagram that shows the different MEG level can be found in [MEF 30.1].

⁵⁸ Revised per L2VPN-N-14.1135-5 on 3/27/14 by JB.

Operator and EVC MEG Level




Figure 8–1 - Operator and EVC MEG Level

8.2.2 CMTS

A CMTS **MUST** support at least one MEG Intermediate Point (MIP) instance per L2VPN [802.1Q]. A CMTS **MAY** support more than one simultaneous MIP instance per L2VPN.

The CMTS **MUST** be configurable per MIP instance with any valid MEG level value (0..7). The CMTS configures a system wide default through mechanisms not specified in this document. See [MEF 30.1] section 7.1 for default values for each MEG level. A CMTS **MUST** support independent configuration of each MIP. Service OAM traffic uses the downstream user priority set for the L2VPN service flow.

A CMTS **MUST** support at least one MEG End Point (MEP) instance per L2VPN [802.1Q].

A CMTS **MAY** support more than one simultaneous MEP instance per L2VPN. This allows CMTS to be configured as a MEP.

8.3 Fault Management⁵⁹

The CM supports a subset of the Fault Management functions defined in [MEF 30.1], [Y.1731] and [802.1Q]. The specific messages supported by the CM are continuity checks, loopback and linktrace. These messages are intended to help with service activation and pro-active management of services.

8.3.1 Continuity Check Messages (CCM)

Continuity Check Messages are used to detect loss of connectivity between two MEPs by transmitting periodic keepalive messages. Continuity Check Messages are sent to all configured remote MEP peers. The CM is expected to support the CCM messages and processes as defined in in Clause 20.1 and 21.6 of [802.1Q]. The CCM frames are sent on the same service flow as that of the L2VPN traffic.

Figure 8–2 provides an overview of how CCM is utilized in the MSO Network. The CM sends and receives CCM frames. When a CMTS is a controller MEP, the CMTS sends and receives CCM frames.

⁵⁹ Revised per L2VPN-N-14.1135-5 on 3/29/14 by JB.

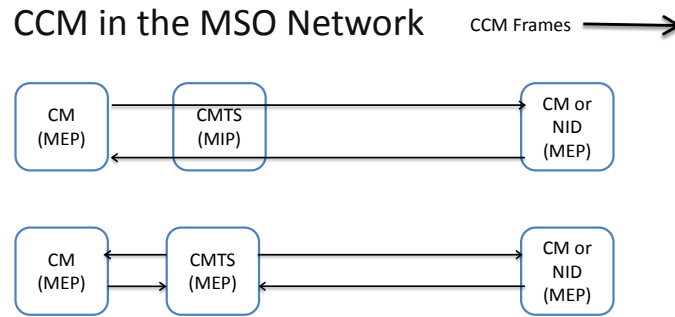
**Figure 8–2 - CCM in the MSO Network**

Table 8–3 describes the flags format of the CCM packet. The Flags format of the CCM packet includes the Remote Defect Indication (RDI) bit and the Transmission period. The Transmission Period is required for CCM transmission and encoded in the 3 least significant bits of the Flag field. Setting the Period field to the value of 1 0 0 indicates a Transmission Period of 1 second. Setting the Period field to the value of 1 0 1 indicates transmission period of 10 seconds. Setting the RDI bit to 1 indicates a defect, while a value of 0 indicates no defect.

Table 8–3 - Flags Format in CCM PDU

8	7	6	5	4	3	2	1
RDI	Reserved (0)				Period		

8.3.1.1 CCM Transmission

The CM and CMTS are expected to support CCM frame transmission periods of both 1 and 10 seconds. Due to the Request/Grant latency of the DOCSIS network, transmission periods less than 1 second are avoided.

The CM is configured with a MAID/MEG ID, MEP ID, List of Peer MEP ID(s), MEG Level, and Transmission Period.

A CM **MUST** support the CCM frame transmission periods of both 1 and 10 seconds. The default value of the CCM frame transmission period is 10 seconds. The CM **MUST NOT** support any frame transmission period of less than 1 second.

The CMTS **MUST** count the number of CCM frames it transmitted to the remote MEP with *dot1agCfmMepCciSentCcms* [802.1Q].

The CMTS is configured with a MAID/MEG ID, MEP ID, List of Peer MEP ID(s), MEG Level, and Transmission Period.

The CMTS **MUST** support the CCM frame transmission periods of both 1 and 10 seconds. The CMTS **MUST NOT** support any frame transmission period of less than 1 second. The CMTS **MUST** count the number of CCM frames it transmitted to the remote MEP with *dot1agCfmMepCciSentCcms* [802.1Q].

8.3.1.2 CCM Reception

The CM **MUST** count the number of CCM frames it received.

The CM takes action when CCM frames are not received from peer MEP within 3.5 times the transmission period. In case of connectivity failure, the CM **MUST** change Remote Defect Indication (RDI) bit in flag format from 0 to 1.

The CM MUST:

- log an event in the CM's local log as defined in Annex D of [OSSIV3.0],
- generate a syslog message if syslog notification is enabled as defined in Annex D of [OSSIV3.0],
- generate an SNMP notification alarm using the dot1agCfmFaultAlarm [802.1Q] if SNMP traps are enabled.

The CM MUST support the DefRDICCM and DefRemoteCCM defects [802.1Q]. The CM SHOULD support the DefMACstatus, DefErrorCCM, and DefXconCCM defects [802.1Q].

The CMTS MUST count the number of CCM frames it received from peer MEP.

The CMTS takes action when CCM frames are not received from peer MEP within 3.5 times the transmission period. In case of connectivity failure, the CMTS MUST change RDI bit in the flag format from 0 to 1.

The CMTS MUST:

- log an event in the CMTS's local log as defined in Annex D of [OSSIV3.0],
- generate a syslog message if syslog notification is enabled as defined in Annex D of [OSSIV3.0],
- generate an SNMP notification alarm using the dot1agCfmFaultAlarm [802.1Q] if SNMP traps are enabled.

The CMTS MUST support the DefRDICCM and DefRemoteCCM defects [802.1Q]. The CMTS SHOULD support the DefMACstatus, DefErrorCCM, and DefXconCCM defects [802.1Q].

8.3.2 Loopback

Loopback messages verify connectivity by transmitting OAM "pings" and receiving the responses. The requirements described here are considered 'lightweight' and a full description of Loopback functions and PDU format can be found on Clause 7.2, 9.3 and 9.4 of [Y.1731] and Clause 20.2 and 21.7 of [802.1Q].

Figure 8–3 below shows an overview of Loopback functionality in the MSO Network. The CM is a responder MEP and responds to loopback frames. When a CMTS is a Controller MEP, the CMTS initiates Loopback frames toward the CM.

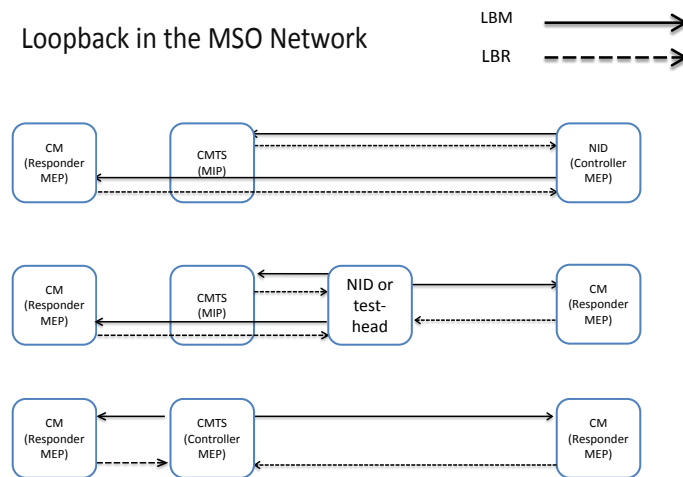


Figure 8–3 - Loopback in the MSO Network

8.3.2.1 Loopback Transmission

The CM has no requirements for transmitting Loopback Message (LBM). The CM functions only as a responder MEP.

The CMTS MUST transmit unicast LBM frame toward remote the MEP or MIP. The CMTS MUST be able to be configured with a valid MEG Level (0..7). The CMTS MUST configure a valid unicast MAC address of the remote MEP or MIP to which the LBM frame is intended.

Unicast LBM frames are transmitted on an on-demand basis.

8.3.2.2 Loopback Reception

The CM MUST respond to a unicast LBM with a unicast Loopback Message Reply (LBR) and therefore implement the Loopback Responder function. The CMTS MUST respond to a unicast LBM with a unicast LBR and, therefore, implement the Loopback Responder function.

Every field in the Unicast LBM frame is copied to the LBR frame except the following:

- Source and Destination MAC addresses are swapped
- OpCode field is changed from 3 to 2 which indicates the frame is now LBR

8.3.3 Linktrace

Linktrace messages identify the path through a network by transmitting OAM frames and receiving responses from various points along the transmission path. The requirements described here are considered 'lightweight' and full description of Linktrace functions and PDU format can be found on Clause 7.3, 9.5 and 9.6 of [Y.1731] and Clause 20.3, 21.8 and 21.9 of [802.1Q].

Figure 8–4 below shows an overview of Linktrace in the MSO Network. The CM is a responder MEP and responds to linktrace frames. When a CMTS is a Controller MEP, the CMTS initiates linktrace frames toward the CM.




Figure 8–4 - Linktrace functionality in the MSO Network

8.3.3.1 Linktrace Transmission

The CM has no requirements for transmitting Linktrace Message (LTM). The CM functions only as a responder MEP.

The CMTS MUST forward the LTM frame that it receives from a MEP toward the next hop (i.e., CM). The CMTS MUST decrement the TTL field on the LTM frame by 1 when transmitting toward the next hop.

The CMTS MUST transmit the LTM frame toward the remote MEP or MIP with a valid MEG Level (0..7). The CMTS MUST include a target MAC address (the MAC address of MEP or MIP). The CMTS MUST be configured with correct value in the TTL field.

A CMTS MUST support Linktrace status and results using the 'TransmitLtmResult', 'LtmTransmitted' and 'LtrReceived' objects from the dot1agCfmMepTable [802.1Q] and mefSoamLtStatsTable [MEF 31].

8.3.3.2 Linktrace Reception

The CM MUST respond to a unicast LTM with a unicast Linktrace Reply (LTR) and, therefore, implement the Linktrace Responder function. The CMTS MUST respond to a LTM with a LTR and, therefore, implement the Linktrace Responder function. Every field in the LTM frame is copied to the LTR except the following:

- Source and Destination MAC addresses are swapped.
- OpCode field is changed from 5 to 4, which indicates the frame is now LTR.

A CM MUST support Linktrace status and results using the 'TransmitLtmResult', 'LtmTransmitted' and 'LtrReceived' objects from the dot1agCfmMepTable [802.1Q] and mefSoamLtStatsTable [MEF 31]. A CMTS MUST support Linktrace status and results using the 'TransmitLtmResult', 'LtmTransmitted' and 'LtrReceived' objects from the dot1agCfmMepTable [802.1Q] and mefSoamLtStatsTable [MEF 31].

8.4 Performance Management⁶⁰

Performance Management functions are used to measure the performance, quality and availability of an active service. There are many different ways to measure performance but this specification will focus on only two mechanisms: Frame Delay and Frame Loss.

Additionally, a CM MAY support the performance benchmark test suites as defined in [RFC 2544] or [Y.1564].

8.4.1 Frame Delay

Frame Delay calculation can be done in a configuration which is dual-ended or single-ended. Dual-ended requires configuration on both MEPs and clock synchronization. Single-ended delay PM requires configuration on one side of the network to initiate Frame Delay Messages and receive Frame Delay Reply from far end and calculates the delay. Clock synchronization is not required for single-ended delay PM functions.

One-way Frame Delay measurement can be calculated with dual-ended PM functions and two-way Frame Delay measurement can be calculated with single-ended PM function.

The requirements described here are considered 'lightweight' and full description of Frame Delay and PDU format can be found on clause 8.2, 9.15 and 9.16 of [Y.1731].

Figure 8–5 below show an overview of Frame Delay in the MSO Network. The CM is a responder MEP and responds to frame delay frames. When a CMTS is a Controller MEP, the CMTS initiates frame delay frames toward the CM. The CMTS also calculates the round-trip delay using its own timestamp.

⁶⁰ Revised per L2VPN-N-14.1135-5 on 3/31/14 by JB.

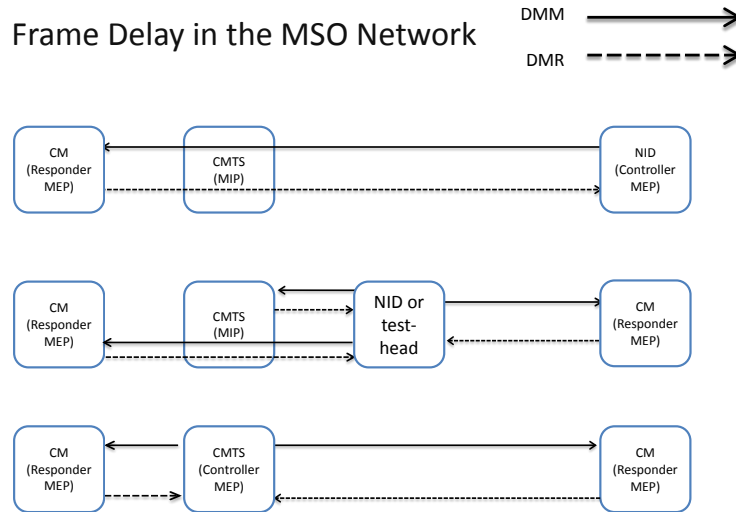


Figure 8–5 - Frame Delay functionality in the MSO Network

8.4.1.1 Frame Delay Transmission

The CM has no requirements for initiating the Frame Delay Measurement Message (DMM). The CM functions only as a responder MEP.

The CMTS MUST transmit a unicast DMM frame toward the remote MEP with TxTimeStampf value. The CMTS MUST be configured to a valid MEG level (0..7). The CMTS will need to include a valid MAC address in the unicast DMM frame.

8.4.1.2 Frame Delay Reception

The CM MUST support Frame Delay using single-ended delay PM functions and therefore implement the Delay Measurement Responder function. The CM will act as a responder to the DMM frame and all Frame Delay measurement is calculated at the other MEP.

A CMTS or a NID (as in Figure 8–4) will initiate a unicast DMM frame towards the CM. The CM MUST respond to the unicast DMM frame with a unicast Delay Measurement Reply (DMR) frame.

When a CM responds to the DMM with a DMR, every field in the DMM frame is copied to the DMR frame except:

- Source and Destination MAC addresses are swapped
- OpCode is changed from 47 to 46, which indicates the frame is now DMR

A CM MUST support number of Frame Delay PDUs it received and sent using the 'mefSoamDmCurrentStatsSoamPdusReceived' and 'mefSoamDmCurrentStatsSoamPdusSent' from [MEF 36].

A CMTS MUST support number of Frame Delay PDUs it received and sent using the 'mefSoamDmCurrentStatsSoamPdusReceived' and 'mefSoamDmCurrentStatsSoamPdusSent' from [MEF 36].

8.4.1.3 Frame Delay Calculation

Upon receiving a DMR frame, the CMTS or a NID will calculate round-trip delay based on the time stamp values. The full calculation of frame delay can be found on Clause 8.2.2.3 [Y.1731].

8.4.2 Synthetic Frame Loss

Synthetic Frame Loss measures frame loss in a network using synthetic frames instead of customer's data frames. The loss is calculated by a MEP by counting the number of synthetic frames it sends and receives.

The requirements described here are considered ‘lightweight’ and full description of Frame Delay and PDU format can be found on clause 8.4, 9.22 and 9.23 of [Y.1731].

Figure 8–6 below shows an overview of Frame Loss Functionality in the MSO Network. The CM is a responder MEP and responds to frame loss frames. When a CMTS is a Controller MEP, the CMTS initiates frame loss frames toward the CM.




Figure 8–6 - Frame Loss Functionality in the MSO Network

8.4.2.1 Synthetic Frame Loss Transmission

The CM has no requirements for initiating Synthetic Loss Message (SLM). The CM functions only as a responder MEP.

The CMTS MUST transmit a unicast SLM frame with a valid MEG level (0..7). The CMTS MUST include the destination MAC address (the address of the remote MEP) in the unicast SLM frame. The CMTS MUST configure a Test ID in the unicast SLM frame.

8.4.2.2 Synthetic Frame Loss Reception

The CM is expected to support Frame Loss using single-ended loss PM function and therefore implement the Loss Measurement Responder function. The CM MUST respond to the SLM frame with a Synthetic Loss Reply (SLR).

When CM responds to the SLM with an SLR, every field in SLM frame is copied to the SLR frame except:

- Source and Destination MAC addresses are swapped,
- OpCode is changed from 55 to 54, which indicates the frame is now SLR,
- Responder MEP ID is changed to MEP’s own identity in the MEG,
- TxFCb: the value of the local counter RxFC1 at the time of transmission.

A CM MUST support reporting the number of Frame Loss PDUs it received and sent using the ‘mefSoamLmCurrentStatsSoamPdusSent’ and ‘mefSoamLmCurrentStatsSoamPdusReceived’ [MEF 36].

8.4.2.3 Synthetic Frame Loss Calculation

Upon receiving an SLR frame, the CMTS or a NID will calculate frame loss for a given measurement period. The full calculation of synthetic frame loss can be found on Clause 8.4.1.3 of [Y.1731].

9 LAYER 2 CONTROL PROTOCOL HANDLING ⁶¹

Layer 2 Control Protocols (L2CP) are a collection of several Layer 2 protocols which are used for various Ethernet control purposes (e.g., Spanning Tree Protocol). Some Layer 2 Control protocols share the same destination MAC address and are identified by additional fields such as the Ethertype and a protocol identifier. Identification of service frames carrying Layer 2 Control Protocols is specified based on destination MAC address.

Per [MEF 6.1.1], L2CP protocols are configured to 'tunnel', 'peer', or 'discard'. 'Discard' means that the network will filter ingress L2CP frames. In this specification, such behavior will be referred to as 'filter'. 'Peer' means that network equipment will actively participate with the protocol. L2CP 'peering' is not required across a DOCSIS network. 'Tunnel' means that Service Frames containing the protocol will be forwarded across the network to the destination without any additional processing. In this specification, such behavior will be referred to as 'forward'.

The CM MUST filter or forward upstream L2CP frames per Table 9–1 for each L2VPN unless overridden by operator configuration such as DOCSIS UDCs.

The CMTS MUST filter or forward downstream L2CP frames per Table 9–1 for each L2VPN unless overridden by operator configuration on the CMTS.

The CM is expected to forward all downstream L2CP frames that it receives as a part of the L2VPN. Likewise, the CMTS forwards all upstream L2CP frames that it receives as a part of the L2VPN.

Table 9–1 - L2CP Frames and Actions

Destination MAC Address L2CP Frames	L2CP Action
01-80-C2-00-00-00	Forward
01-80-C2-00-00-01 through 01-80-C2-00-00-0A	Filter
01-80-C2-00-00-0B	Forward
01-80-C2-00-00-0C	Forward
01-80-C2-00-00-0D	Forward
01-80-C2-00-00-0E	Filter
01-80-C2-00-00-0F	Forward
01-80-C2-00-00-20 through 01-80-C2-00-00-2F	Forward
01-80-C2-00-00-30 through 01-80-C2-00-00-3F	Forward Note: If SOAM is configured, the CM may need to process multicast messages for its configured MEG level.

⁶¹ New section added per L2VPN-N-12.1065-11 on 9/19/12 by CP.

Annex A CMTS DOCS-L2VPN-MIB Requirements (Normative)

A CMTS MUST implement the DOCS-L2VPN-MIB. A CM does not implement DOCS-L2VPN-MIB.

A.1 DOCS-L2VPN-MIB Conformance

Legend:

M	Mandatory
NA	Not Applicable
RO	Read-Only
RC	Read-Create

DOCS-L2VPN-MIB ⁶²				
docsL2vpnIdToIndexTable (Point-to-Point and Multipoint)				
Object	CM	Access	CMTS	Access
docsL2vpnIdToIndexIdx	NA	NA	M	RO
docsL2vpnIndexToIdTable (Point-to Point and Multipoint)				
Object	CM	Access	CMTS	Access
docsL2vpnIndexToIdId	NA	NA	M	RO
docsL2vpnCmTable				
docsL2vpnCmCompliantCapability	NA	NA	M	RO
docsL2vpnCmDutFilteringCapability	NA	NA	M	RO
docsL2vpnCmDutCMIM	NA	NA	M	RO
docsL2vpnCmDhcpSnooping	NA	NA	M	RO

⁶² New MIB objects added per L2VPN-N-12.1065-11 on 9/19/12 by CP.

DOCS-L2VPN-MIB ⁶²				
docsL2vpnVpnCmTable				
Object	CM	Access	CMTS	Access
docsL2vpnVpnCmDhcpSnooping	NA	NA	M	RO
docsL2vpnVpnCmCMIM	NA	NA	M	RO
docsL2vpnVpnCmVendorSpecific	NA	NA	M	RO
docsL2vpnVpnCmStatsTable (Point-to-Point and Multipoint)				
Object	CM	Access	CMTS	Access
docsL2vpnVpnCmStatsUpstreamPkts	NA	NA	M	RO
docsL2vpnVpnCmStatsUpstreamBytes	NA	NA	M	RO
docsL2vpnVpnCmStatsUpstreamDiscards	NA	NA	M	RO
docsL2vpnVpnCmStatsDownstreamPkts	NA	NA	M	RO
docsL2vpnVpnCmStatsDownstreamBytes	NA	NA	M	RO
docsL2vpnVpnCmStatsDownstreamDiscards	NA	NA	M	RO
docsL2vpnPortStatusTable (Point-to-Point and Multipoint)				
Object	CM	Access	CMTS	Access
docsL2vpnPortStatusSAID	NA	NA	M	RO
docsL2vpnSfStatusTable (Point-to-Point and Multipoint)				
Object	CM	Access	CMTS	Access
docsL2vpnSfStatusL2vpnId	NA	NA	M	RO

DOCS-L2VPN-MIB ⁶²				
docsL2vpnSfStatusUpstreamUserPriority	NA	NA	M	RO
docsL2vpnSfStatusVendorSpecific	NA	NA	M	RO
docsL2vpnPktClassTable (Point-to-Point and Multipoint)				
Object	CM	Access	CMTS	Access
docsL2vpnPktClassL2vpnIdx ⁶³	NA	NA	M	RO
docsL2vpnPktClassUserPriRangeLow	NA	NA	M	RO
docsL2vpnPktClassUserPriRangeHigh	NA	NA	M	RO
docsL2vpnPktClassCmim	NA	NA	M	RO
docsL2vpnPktClassVendorSpecific	NA	NA	M	RO
docsL2vpnCmNsiTable (Point-to-Point Only)				
Object	CM	Access	CMTS	Access
docsL2vpnCmNsiEncapSubtype	NA	NA	M	RO
docsL2vpnCmNsiEncapValue,	NA	NA	M	RO
docsL2vpnCmNsiAGI,	NA	NA	M	RO
docsL2vpnCmNsiSAII	NA	NA	M	RO
docsL2vpnCmNsiMode	NA	NA	M	RO
docsL2vpnCmNsiTpidTransUS	NA	NA	M	RO
docsL2vpnCmNsiTpidTransDS	NA	NA	M	RO
docsL2vpnCmVpnCpeTable (Multipoint Only)				

⁶³ Modified per L2VPN-N-06.0315-2 by PO 11/22/06.

DOCS-L2VPN-MIB ⁶²				
Object	CM	Access	CMTS	Access
docsL2vpnCmVpnCpeMacAddress	NA	NA	M	RO
docsL2vpnVpnCmCpeTable (Multipoint only)				
Objects	CM	Access	CMTS	Access
docsL2vpnVpnCmCpeMacAddress	NA	NA	M	RO
docsL2vpnDot1qTpFdbExtTable (Multipoint only)				
Objects	CM	Access	CMTS	Access
docsL2vpnDot1qTpFdbExtTransmitPkts	NA	NA	M	RO
docsL2vpnDot1qTpFdbExtReceivePkts	NA	NA	M	RO
docsL2vpnDot1qTpGroupExtTable (Multipoint only)				
Objects	CM	Access	CMTS	Access
docsL2vpnDot1qTpGroupExtTransmitPkts	NA	NA	M	RO
docsL2vpnDot1qTpGroupExtReceivePkts	NA	NA	M	RO
docsL2vpnVpnSfStatsTable (Point-to-Point and Multipoint)				
Object	CM	Access	CMTS	Access
docsL2vpnVpnSfStatsUpstreamPkts	NA	NA	M	RO
docsL2vpnVpnSfStatsUpstreamBytes	NA	NA	M	RO
docsL2vpnVpnSfStatsUpstreamDiscards	NA	NA	M	RO
docsL2vpnVpnSfStatsDownstreamPkts	NA	NA	M	RO

DOCS-L2VPN-MIB ⁶²				
docsL2vpnVpnSfStatsDownstreamBytes	NA	NA	M	RO
docsL2vpnVpnSfStatsDownstreamDiscards	NA	NA	M	RO

A.2 DOCS-L2VPN-MIB Definitions ⁶⁴

DOCS-L2VPN-MIB DEFINITIONS ::= BEGIN

IMPORTS

```

    MODULE-IDENTITY,
    OBJECT-TYPE,
    Unsigned32,
    Integer32,
    Counter32
        FROM SNMPv2-SMI

    TEXTUAL-CONVENTION,
    TruthValue,
    MacAddress
        FROM SNMPv2-TC

    MODULE-COMPLIANCE,
    OBJECT-GROUP
        FROM SNMPv2-CONF

    ifIndex
        FROM IF-MIB

    dot1dBasePort
        FROM BRIDGE-MIB

    dot1qFdbId,
    dot1qTpFdbAddress,
    dot1qVlanIndex,
    dot1qTpGroupAddress
        FROM Q-BRIDGE-MIB

    docsIfCmtsCmStatusIndex
        FROM DOCS-IF-MIB

    docsQosServiceFlowId,
    docsQosPktClassId
        FROM DOCS-QOS-MIB

    clabProjDocsis,
    DocsL2vpnIfList
        FROM CLAB-DEF-MIB;

```

docsL2vpnMIB MODULE-IDENTITY

```

    LAST-UPDATED "201409240000Z" -- September 24, 2014.
    ORGANIZATION "Cable Television Laboratories, Inc."
    CONTACT-INFO
        "Postal: Cable Television Laboratories, Inc.
        858 Coal Creek Circle
        Louisville, Colorado 80027-9750
        U.S.A.
        Phone: +1 303-661-9100

```

⁶⁴ MIB definition changed per L2VPN-N-06.0316-2 and L2VPN-N-06.0315-2 by PO 11/22/06. Entire section replaced per L2VPN-N-07.0571 on 1/24/08 by KN. Further revised per L2VPN-N-12.1065-11 on 9/19/12 by CP, section replaced per L2VPN-N-13.1109-3 on 7/23/13 by PO. Section modified per L2VPN-N-14.1186-1 on 1/15/15 by PO.

```

    Fax:      +1 303-661-9199
    E-mail: mibs@cablelabs.com"
DESCRIPTION
    "This is the management MIB for devices complying to the
    DOCSIS L2VPN Feature.
    Copyright 2012 Cable Television Laboratories, Inc.
    All rights reserved."
REVISION "201503050000Z =" - March 5, 2015
DESCRIPTION
    "Revised Version includes ECN L2VPN-N-14.1186-1
    and published as I14"
REVISION "201308080000Z" -- August 8, 2013
DESCRIPTION
    "Revised Version includes ECN L2VPN-N-13.1109-3
    and published as I11"
REVISION "201210040000Z" -- October 04, 2012
DESCRIPTION
    "Revised Version includes ECN L2VPN-N-12.1065-11
    and published as I10"
REVISION "200802150000Z" -- February 15, 2008
DESCRIPTION
    "Revised Version includes ECN L2VPN-N-07.0571-1
    and published as I07"
REVISION "200612220000Z" -- December 22, 2006
DESCRIPTION
    "Published as part of the DOCSIS
    Layer 2 Virtual Private Networks Specification
    CM-SP-L2VPN-I03-061222."
REVISION "200607280000Z" -- July 28, 2006
DESCRIPTION
    "Published as part of the DOCSIS
    Layer 2 Virtual Private Networks Specification
    CM-SP-L2VPN-I02-060728."
REVISION "200603280000Z" -- March 28, 2006
DESCRIPTION
    "Initial version, published as part of the DOCSIS
    Layer 2 Virtual Private Networks Specification
    CM-SP-L2VPN-I01-060328."
::= { clabProjDocsis 8 }

-----
--
-- Textual Conventions
--
Docsl2vpnIdentifier ::= TEXTUAL-CONVENTION
    DISPLAY-HINT "255a"
    STATUS         current
    DESCRIPTION
        "An externally administered octet string identifying an
        L2VPN. An implementation MUST support a length of at least
        16 octets. The octet string is used as an index. As such,
        the CMTS enforces that objects of type Docsl2vpnIdentifier
        are unique per CMTS. An MSO is encouraged to define
        Docsl2vpnIdentifier values as globally unique."
    SYNTAX         OCTET STRING (SIZE(1..16))

Docsl2vpnIndex ::= TEXTUAL-CONVENTION
    STATUS         current
    DESCRIPTION
        "An integer value locally generated by the agent for each
        known Docsl2vpnIdentifier administrative identifier. It is
        intended to be used as a short index for tables in this MIB
        module in lieu of an object of the type

```

```

        DocsL2vpnIdentifier."
SYNTAX      Unsigned32 (0..4294967295)

DocsNsiEncapSubtype ::= TEXTUAL-CONVENTION
    STATUS      current
    DESCRIPTION
        "An enumerated integer that defines the default
        encapsulation on NSI ports of an L2VPN-forwarded packet.
        A CMTS implementation MUST support ieee802.1q(2).
        A CMTS MAY omit support for all NSI encapsulations
        other than ieee802.1q(2)."
```

SYNTAX INTEGER {

```

    other(1),
    ieee8021q(2),
    ieee8021ad(3),
    mpls(4),
    l2tpv3(5),
    ieee8021ah(6)
}
```

DocsNsiEncapValue ::= TEXTUAL-CONVENTION

```

    STATUS      current
    DESCRIPTION
        "The encapsulation value for L2VPN forwarded packets on NSI
        ports. The value of an object of this type depends on the
        value of an associated object of type DocsEncapSubtype:

        other(1): vendor specific,
        ieee8021q(2): 802.1Q tag with VLAN ID in lower 12 bits,
        ieee8021ad(3): pair of 16-bit values with service provider
        in lower 12 bits of the first 16-bit value and customer
        VLAN ID in the lower 12 bits of the second 16-bit value,
        mpls(4): has to be zero length string,
        l2tpv3(5): has to be zero length string,
        ieee8021q(6): IEEE 802.1Q encapsulation which consists of B-SA,
        B-DA and 48-bit I-Tag."
```

SYNTAX OCTET STRING

DocsNsiModeSubtype ::= TEXTUAL-CONVENTION

```

    STATUS      current
    DESCRIPTION
        "An enumerated integer that defines L2VPN Mode. Two modes
        are supported: Encapsulation Mode and Transport Mode."
```

SYNTAX INTEGER {

```

    none(1),
    encapsulation(2),
    transport(3)
}
```

```

-- Placeholder for notifications
--
docsL2vpnMIBNotifications OBJECT IDENTIFIER ::= { docsL2vpnMIB 0 }

-- None defined

--
-- L2VPN MIB Objects
--

docsL2vpnMIBObjects OBJECT IDENTIFIER ::= { docsL2vpnMIB 1 }
```

```

-----
--
-- Point-to-Point and Point-to-Multipoint
--
-- The following objects are required for both
-- Point-to-Point and Point-to-Multipoint operation.
--
-----
--
-- L2VPN Identifier to L2VPN Index mapping table
--
docsL2vpnIdToIndexTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnIdToIndexEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "Table indexed by the octet string DocsL2vpnIdentifier that
        provides the local agent's internally assigned docsL2vpnIdx
        value for that DocsL2vpnIdentifier value. The mapping of
        DocsL2vpnIdentifier to docsL2vpnIdx is 1-1. The agent
        has to instantiate a row in both docsL2vpnIndexToIdTable and
        docsL2vpnIdToIndexTable for each known L2VPN Identifier."
    ::= { docsL2vpnMIBObjects 1 }

docsL2vpnIdToIndexEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnIdToIndexEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "Maps a DocsL2vpnIdentifier octet string into the local
        agent's locally assigned docsL2vpnIdx value."
    INDEX { docsL2vpnId }
    ::= { docsL2vpnIdToIndexTable 1 }

DocsL2vpnIdToIndexEntry ::= SEQUENCE
{
    docsL2vpnId          DocsL2vpnIdentifier,
    docsL2vpnIdToIndexIdx DocsL2vpnIndex
}

docsL2vpnId OBJECT-TYPE
    SYNTAX      DocsL2vpnIdentifier
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "An externally configured octet string that identifies an
        L2VPN."
    ::= { docsL2vpnIdToIndexEntry 1 }

docsL2vpnIdToIndexIdx OBJECT-TYPE
    SYNTAX      DocsL2vpnIndex
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "An internally assigned index value for a known L2VPN."
    ::= { docsL2vpnIdToIndexEntry 2 }

-----
--
-- L2VPN Index to L2VPN Identifier mapping tables
--

```

```
docsL2vpnIndexToIdTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnIndexToIdEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "Table indexed by agent's local docsL2vpnIdx that provides
        the global L2VPN Identifier. The mapping of docsL2vpnIdx to
        DocsL2vpnIdentifier is 1-1. The agent has to instantiate a
        row in both docsL2vpnIndexToIdTable and
        docsL2vpnIdToIndexTable for each known L2VPN."
    ::= { docsL2vpnMIBObjects 2 }
```

```
docsL2vpnIndexToIdEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnIndexToIdEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "Provides the L2VPN Identifier for each locally-assigned
        L2vpn Index."
    INDEX { docsL2vpnIdx }
    ::= { docsL2vpnIndexToIdTable 1 }
```

```
DocsL2vpnIndexToIdEntry ::= SEQUENCE
{
    docsL2vpnIdx          DocsL2vpnIndex,
    docsL2vpnIndexToIdId DocsL2vpnIdentifier
}
```

```
docsL2vpnIdx OBJECT-TYPE
    SYNTAX      DocsL2vpnIndex
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "An internally assigned index value for a known L2VPN."
    ::= { docsL2vpnIndexToIdEntry 1 }
```

```
docsL2vpnIndexToIdId OBJECT-TYPE
    SYNTAX      DocsL2vpnIdentifier
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "An administered octet string that externally identifies an
        L2VPN."
    ::= { docsL2vpnIndexToIdEntry 2 }
```

```
-----
--
```

```
-- L2VPN CM Table
-- Point-to-Point and Multipoint mode
--
```

```
docsL2vpnCmTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnCmEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This table describes L2VPN per-CM information that
        is in common with all L2VPNs for the CM, regardless
        of forwarding mode."
    ::= { docsL2vpnMIBObjects 3 }
```

```
docsL2vpnCmEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnCmEntry
    MAX-ACCESS  not-accessible
```

```

STATUS      current
DESCRIPTION
    "An entry is indexed by Cable Modem Index that
    describes L2VPN information for a single CM that is in
    common with all L2VPNs implemented by the CM,
    regardless of the L2VPN forwarding mode.

    An entry in this table is created for every CM that
    registers with a forwarding L2VPN encoding."
INDEX { docsIfCmtsCmStatusIndex }
::= { docsL2vpnCmTable 1 }

DocsL2vpnCmEntry ::= SEQUENCE {
    docsL2vpnCmCompliantCapability      TruthValue,
    docsL2vpnCmDutFilteringCapability   TruthValue,
    docsL2vpnCmDutCMIM                  DocsL2vpnIfList,
    docsL2vpnCmDhcpSnooping             DocsL2vpnIfList
}

docsL2vpnCmCompliantCapability OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "This object reports whether an L2VPN forwarding CM is
        compliant with the DOCSIS L2VPN specification, as reported
        in the L2VPN Capability encoding in the CM's registration
        request message.

        If the capability encoding was omitted, this object has to
        report the value false(2).
        ::= { docsL2vpnCmEntry 1 }

docsL2vpnCmDutFilteringCapability OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "This object reports whether an L2VPN forwarding CM is
        capable of Downstream Unencrypted Traffic (DUT) Filtering,
        as reported in the CM's registration request message.

        If the capability encoding was omitted, this object has to
        report the value false(2).
        ::= { docsL2vpnCmEntry 2 }

docsL2vpnCmDutCMIM OBJECT-TYPE
    SYNTAX      DocsL2vpnIfList
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "This object reports the value configured in a per-CM
        L2VPN Encoding for Downstream Unencrypted Traffic (DUT)
        Cable Modem Interface Mask (CMIM).

        The DUT CMIM is a bit string with a '1' for each bit
        position K for an internal or external CM interface with
        ifIndex K to which the CM permits DUT to be forwarded. A CM
        capable of DUT filtering MUST discard DUT to interfaces
        with a '0' in the DUT CMIM.

        If a CM's top-level registration request L2VPN Encoding
        contained no DUT CMIM subtype, this object is reported

```

with its default value of a '1' in bit position 0 (corresponding to the eCM's own 'self' host) and a '1' in each bit position K for which an eSAFE interface exists at ifIndex K. In other words, the default DUT CMIM includes the eCM and all eSAFE interfaces.

This value is reported independently of whether the CM is actually capable of performing DUT filtering."

```
::= { docsL2vpnCmEntry 3 }
```

docsL2vpnCmDhcpSnooping OBJECT-TYPE

SYNTAX DocsL2vpnIfList

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"This object reports the value of the Enable DHCP Snooping subtype of a top-level L2VPN Encoding.

It has the syntax of a CM Interface List bitmask and represents a set of CM MAC bridge interfaces corresponding to eSAFE hosts for which the CMTS is enabled to snoop DHCP traffic in order to learn the eSAFE host MAC address on that interface.

Only bits corresponding to eSAFE host MAC addresses may be validly set in this object, including cpe(1) for eRouter and the eSAFE interfaces in bits positions 16 through 31."

```
::= { docsL2vpnCmEntry 4 }
```

```
-----
--
-- L2VPN/CM Table
-- Point-to-Point and Multipoint mode
--
```

docsL2vpnVpnCmTable OBJECT-TYPE

SYNTAX SEQUENCE OF DocsL2vpnVpnCmEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"This table describes the operation of L2VPN forwarding on each CM."

```
::= { docsL2vpnMIBObjects 4 }
```

docsL2vpnVpnCmEntry OBJECT-TYPE

SYNTAX DocsL2vpnVpnCmEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"An entry is indexed by VPN ID and Cable Modem Index that describes the operation of L2VPN forwarding for a single L2VPN on a single CM."

INDEX { docsL2vpnIdx, docsIfCmtsCmStatusIndex }

```
::= { docsL2vpnVpnCmTable 1 }
```

DocsL2vpnVpnCmEntry ::= SEQUENCE {

docsL2vpnVpnCmCMIM DocsL2vpnIfList,

docsL2vpnVpnCmIndividualSAId Integer32,

docsL2vpnVpnCmVendorSpecific OCTET STRING

}

docsL2vpnVpnCmCMIM OBJECT-TYPE

SYNTAX DocsL2vpnIfList

```

MAX-ACCESS    read-only
STATUS        current
DESCRIPTION
    "A Cable Modem Interface Mask represents a set of
    MAC bridge interfaces within the CM. This object
    represents the CMIM within a forwarding per-SF L2VPN
    encoding, which specifies a set of CM MAC bridge
    interfaces to which L2VPN forwarding is restricted.

    If the CMIM Subtype is omitted from a forwarding
    per-SF encoding, its default value includes only
    cpePrimary(1) and cableMac(2), which can be encoded
    with a single octet with the value 0x60."
 ::= { docsL2vpnVpnCmEntry 1 }

docsL2vpnVpnCmIndividualSAId OBJECT-TYPE
    SYNTAX      Integer32 (0..16383)
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The BPI+ Security Association ID in which traffic intended
        for point-to-point forwarding through an individual CM is
        forwarded.

        If the CMTS does not allocate an individual SAID for
        multipoint forwarding (as is recommended), it MUST
        report this object as zero."
 ::= { docsL2vpnVpnCmEntry 2 }

docsL2vpnVpnCmVendorSpecific OBJECT-TYPE
    SYNTAX      OCTET STRING
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object encodes the concatenation of all Vendor
        Specific Subtype encodings that appeared in any
        registration per-CM L2VPN Encoding associated with this
        entry."
 ::= { docsL2vpnVpnCmEntry 3 }

-----
--
--
-- L2VPN/CM Statistics Table
--   Point-to-Point and Multipoint mode
--
docsL2vpnVpnCmStatsTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnVpnCmStatsEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This table contains statistics for forwarding of
        packets to and from a CM on each VPN."
 ::= { docsL2vpnMIBObjects 5 }

docsL2vpnVpnCmStatsEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnVpnCmStatsEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "An entry is indexed by VPN ID and Cable Modem Index."
    INDEX { docsL2vpnIdx, docsIfCmtsCmStatusIndex }
 ::= { docsL2vpnVpnCmStatsTable 1 }

```

```

DocsL2vpnVpnCmStatsEntry ::= SEQUENCE {
    docsL2vpnVpnCmStatsUpstreamPkts      Counter32,
    docsL2vpnVpnCmStatsUpstreamBytes     Counter32,
    docsL2vpnVpnCmStatsUpstreamDiscards  Counter32,
    docsL2vpnVpnCmStatsDownstreamPkts    Counter32,
    docsL2vpnVpnCmStatsDownstreamBytes   Counter32,
    docsL2vpnVpnCmStatsDownstreamDiscards Counter32
}

docsL2vpnVpnCmStatsUpstreamPkts OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The number of L2vpn-forwarded packets received
        from this instance's Cable Modem on
        this instance's L2VPN."
    ::= { docsL2vpnVpnCmStatsEntry 1 }

docsL2vpnVpnCmStatsUpstreamBytes OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The number of L2vpn-forwarded bytes received
        from this instance's Cable Modem on
        this instance's L2VPN."
    ::= { docsL2vpnVpnCmStatsEntry 2 }

docsL2vpnVpnCmStatsUpstreamDiscards OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The number of L2-forwarded packets
        discarded from this instance's
        Cable Modem on this instance's VPN."
    ::= { docsL2vpnVpnCmStatsEntry 3 }

docsL2vpnVpnCmStatsDownstreamPkts OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The number of L2-forwarded packets
        transmitted to this instance's
        Cable Modem on this instance's VPN."
    ::= { docsL2vpnVpnCmStatsEntry 4 }

docsL2vpnVpnCmStatsDownstreamBytes OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The number of L2-forwarded bytes
        transmitted to this instance's
        Cable Modem on this instance's VPN."
    ::= { docsL2vpnVpnCmStatsEntry 5 }

docsL2vpnVpnCmStatsDownstreamDiscards OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS       current

```

```

DESCRIPTION
    "The number of L2-forwarded packets that were discarded
    before they could be transmitted to this instance's
    Cable Modem on this instance's VPN."
::= { docsL2vpnVpnCmStatsEntry 6 }

-----
--
--
-- VPN Port Status Table
-- (Point-to-Point and Multipoint mode)
--
docsL2vpnPortStatusTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnPortStatusEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This table displays summary information for the
        run-time state of each VPN that is currently operating
        on each bridge port."
    ::= { docsL2vpnMIBObjects 6 }

docsL2vpnPortStatusEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnPortStatusEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "Information specific to the operation of L2VPN forwarding
        on a particular CMTS 'bridge port'. A CMTS 'bridge port'
        may be defined by the CMTS vendor, but is advantageously a
        single DOCSIS MAC Domain."
    INDEX { dot1dBasePort, docsL2vpnIdx }
    ::= { docsL2vpnPortStatusTable 1 }

DocsL2vpnPortStatusEntry ::= SEQUENCE {
    docsL2vpnPortStatusGroupSAID  Integer32
}

docsL2vpnPortStatusGroupSAID OBJECT-TYPE
    SYNTAX      Integer32 (0..16383)
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The Group SAID associated with this VPN on a
        particular CMTS MAC domain. This SAID is used to encrypt
        all downstream flooded bridge traffic sent to CMs on
        this VPN and CMTS MAC domain bridge port.

        A value of '0' means there is no associated Group SAID for
        this VPN and bridge port, e.g., if the L2VPN uses
        point-to-point individual SAIDs only for forwarding.

        A bridge port that is not a CMTS MAC
        domain will report a value of '0'."
    ::= { docsL2vpnPortStatusEntry 1 }

-----
--
--
-- L2VPN Service Flow Status Table
-- (Point-to-Point and Multipoint mode)
--
-- This table has a row for each upstream SF with a per-SF L2VPN
-- Encoding.
--

```

```

docsL2vpnSfStatusTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnSfStatusEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This table displays SF-specific L2VPN forwarding status
        for each upstream service flow configured with a per-SF
        L2VPN Encoding.

        Objects which were signaled in a per-SF L2VPN Encoding but
        apply for the entire CM are shown in the
        docsL2vpnVpnCmTable."
    ::= { docsL2vpnMIBObjects 7 }

docsL2vpnSfStatusEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnSfStatusEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "SF-specific L2VPN forwarding status information for each
        upstream service flow configured with a per-SF L2VPN
        Encoding. The ifIndex is of type docsCableMacLayer(127)."
    INDEX { ifIndex, docsQosServiceFlowId }
    ::= { docsL2vpnSfStatusTable 1 }

DocsL2vpnSfStatusEntry ::= SEQUENCE {
    docsL2vpnSfStatusL2vpnId          OCTET STRING,
    docsL2vpnSfStatusUpstreamUserPriority  Unsigned32,
    docsL2vpnSfStatusVendorSpecific    OCTET STRING
}

docsL2vpnSfStatusL2vpnId OBJECT-TYPE
    SYNTAX      OCTET STRING
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object represents the value of the L2VPN Identifier
        subtype of a per-SF L2VPN Encoding."
    ::= { docsL2vpnSfStatusEntry 1 }

docsL2vpnSfStatusUpstreamUserPriority OBJECT-TYPE
    SYNTAX      Unsigned32 (0..7)
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object provides the configured Ingress User Priority
        subtype of a per-SF L2VPN Encoding for this CM. If the
        subtype was omitted, this object's value is zero."
    ::= { docsL2vpnSfStatusEntry 2 }

docsL2vpnSfStatusVendorSpecific OBJECT-TYPE
    SYNTAX      OCTET STRING
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object provides the set of configured Vendor Specific
        subtypes within a per-SF L2VPN Encoding for a CM. If no
        Vendor Specific subtype was specified, this object is a
        zero length octet string. If one or more Vendor Specific
        subtype parameters was specified, this object represents
        the concatenation of all such subtypes."
    ::= { docsL2vpnSfStatusEntry 3 }

```

```

-----
--
-- L2VPN Classifier Table
-- (Point-to-Point and Multipoint mode)
--

docsL2vpnPktClassTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnPktClassEntry
    MAX-ACCESS   not-accessible
    STATUS      current
    DESCRIPTION
        "This table provides the L2VPN-specific objects for
        packet classifiers that apply to only L2VPN traffic.
        The indices of this table are a subset of the
        indices of classifiers in docsQosPktClassTable."
    ::= { docsL2vpnMIBObjects 8 }

docsL2vpnPktClassEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnPktClassEntry
    MAX-ACCESS   not-accessible
    STATUS      current
    DESCRIPTION
        "An entry in this table extends a single row
        of docsQosPktClassTable for a rule that applies only to
        downstream L2VPN forwarded packets.
        The index ifIndex is an ifType of docsCableMaclayer(127)."
```

INDEX {	
ifIndex,	
docsQosServiceFlowId,	
docsQosPktClassId	
}	
::= { docsL2vpnPktClassTable 1 }	

```

DocsL2vpnPktClassEntry ::= SEQUENCE {
    docsL2vpnPktClassL2vpnIdx      DocsL2vpnIndex,
    docsL2vpnPktClassUserPriRangeLow Unsigned32,
    docsL2vpnPktClassUserPriRangeHigh Unsigned32,
    docsL2vpnPktClassCMIM          DocsL2vpnIfList,
    docsL2vpnPktClassVendorSpecific OCTET STRING
}

docsL2vpnPktClassL2vpnIdx OBJECT-TYPE
    SYNTAX      DocsL2vpnIndex
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The locally assigned L2VPN index corresponding to the VPN
        Identifier subtype of a Downstream Classifier L2VPN
        Encoding."
    ::= { docsL2vpnPktClassEntry 1 }

docsL2vpnPktClassUserPriRangeLow OBJECT-TYPE
    SYNTAX      Unsigned32 (0..7)
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The lower priority of the user Priority Range subtype
        of a Downstream Classifier L2VPN Encoding. If the subtype
        was omitted, this object has value 0."
    ::= { docsL2vpnPktClassEntry 2 }

docsL2vpnPktClassUserPriRangeHigh OBJECT-TYPE

```

```

SYNTAX      Unsigned32 (0..7)
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The higher priority of the user Priority Range subtype
    of a Downstream Classifier L2VPN Encoding. If the subtype
    was omitted, this object has value 7."
 ::= { docsL2vpnPktClassEntry 3 }

```

docsL2vpnPktClassCMIM OBJECT-TYPE

```

SYNTAX      DocsL2vpnIfList
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The Cable Modem Interface Mask (CMIM) signaled in a
    Packet Classifier Encoding. In a Downstream Packet
    Classifier Encoding, a specified CMIM value restricts the
    classifier to match packets with a Destination MAC address
    corresponding to the interfaces indicated in the CMIM mask.
    The eCM self and any eSAFE interface bits correspond to
    the individual eCM and eSAFE host MAC addresses.

    In an Upstream Packet Classifier encoding, a specified CMIM
    value restricts the classifier to match packets with an
    ingress bridge port interface matching the bits in the
    CMIM value.

    If the CMIM subtype was omitted, this object should be
    reported as a zero length octet string."
 ::= { docsL2vpnPktClassEntry 4 }

```

docsL2vpnPktClassVendorSpecific OBJECT-TYPE

```

SYNTAX      OCTET STRING
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "This object provides the set of configured
    Vendor Specific subtypes within a Packet Classifier
    Encoding for a CM. If no Vendor Specific subtype was
    specified, this object is a zero length octet string.
    If one or more Vendor Specific subtype parameters was
    specified, this object represents the concatenation of all
    such subtypes."
 ::= { docsL2vpnPktClassEntry 5 }

```

```

-----
--
-- L2VPN CM NSI Table
-- Point-to-Point Only
--
docsL2vpnCmNsiTable OBJECT-TYPE
SYNTAX      SEQUENCE OF DocsL2vpnCmNsiEntry
MAX-ACCESS  not-accessible
STATUS      current
DESCRIPTION
    "This table describes the NSI configuration for a single
    CM when operating in point-to-point forwarding mode for an
    L2VPN."
 ::= { docsL2vpnMIBObjects 9 }

```

docsL2vpnCmNsiEntry OBJECT-TYPE

```

SYNTAX      DocsL2vpnCmNsiEntry
MAX-ACCESS  not-accessible

```

```

STATUS      current
DESCRIPTION
    "An entry indexed by VPN ID and Cable Modem Index that
    describes the point-to-point forwarding between a single
    NSI encapsulation and a single CM. This table is
    implemented only for a CM forwarding an L2VPN on a
    point-to-point basis. It is associated with a single
    per-CM L2VPN encoding."
INDEX { docsL2vpnIdx, docsIfCmtsCmStatusIndex }
::= { docsL2vpnCmNsiTable 1 }

DocsL2vpnCmNsiEntry ::= SEQUENCE {
    docsL2vpnCmNsiEncapSubtype      DocsNsiEncapSubtype,
    docsL2vpnCmNsiEncapValue        DocsNsiEncapValue,
    docsL2vpnCmNsiAGI               OCTET STRING,
    docsL2vpnCmNsiSAII              OCTET STRING,
    docsL2vpnCmNsiTAII              OCTET STRING,
    docsL2vpnCmNsiMode              DocsNsiModeSubtype,
    docsL2vpnCmNsiTpidTransDS       Integer32,
    docsL2vpnCmNsiTpidTransUS       Integer32
}

docsL2vpnCmNsiEncapSubtype OBJECT-TYPE
    SYNTAX      DocsNsiEncapSubtype
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The General Encapsulation Information (GEI) subtype of the
        Network System Interface (NSI) encapsulation configured
        for the CM."
    ::= { docsL2vpnCmNsiEntry 1 }

docsL2vpnCmNsiEncapValue OBJECT-TYPE
    SYNTAX      DocsNsiEncapValue
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The encapsulation value for L2VPN forwarded packets on NSI
        ports."
    ::= { docsL2vpnCmNsiEntry 2 }

docsL2vpnCmNsiAGI OBJECT-TYPE
    SYNTAX      OCTET STRING
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "This object is the configuration of any Attachment Group
        Identifier subtype in the per-SF L2VPN Encoding
        represented by this row. If the subtype was omitted, this
        object's value is a zero length string."
    ::= { docsL2vpnCmNsiEntry 3 }

docsL2vpnCmNsiSAII OBJECT-TYPE
    SYNTAX      OCTET STRING
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "This object is the configuration of any Source
        Attachment Individual ID subtype in the L2VPN Encoding
        represented by this row. If the subtype was omitted, this
        object's value is a zero length string."
    ::= { docsL2vpnCmNsiEntry 4 }

```

```

docsL2vpnCmNsiTAII OBJECT-TYPE
    SYNTAX      OCTET STRING
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object is the configuration of any Target
        Attachment Individual ID subtype in the L2VPN Encoding
        represented by this row. If the subtype was omitted, this
        object's value is a zero length string."
    ::= { docsL2vpnCmNsiEntry 5 }

docsL2vpnCmNsiMode OBJECT-TYPE
    SYNTAX      DocsNsiModeSubtype
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object indicates L2VPN Mode, which can be Encapsulation Mode
        or Transport Mode.
        This Mode is configured using TLV 43.5.13.
        In the L2VPN Encapsulation Mode IEEE 802.1Q S-VLAN tagging
        or IEEE 802.1Q encapsulation
        is added to upstream frames according to operator provisioning.
        L2VPN Transport Mode is used to forward traffic that belongs
        to a particular EPL service instance without
        adding extra encapsulations or
        tagging or removing existing ones."
    ::= { docsL2vpnCmNsiEntry 6 }

docsL2vpnCmNsiTpidTransDS OBJECT-TYPE
    SYNTAX      Integer32 (0..65535)
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object is the configuration of the TPID used
        in the Downstream Direction (DS) to translate
        the TPID of the frame.
        Value of 0 indicates that no TPID translation is required in the
        Downstream Direction."
    ::= { docsL2vpnCmNsiEntry 7 }

docsL2vpnCmNsiTpidTransUS OBJECT-TYPE
    SYNTAX      Integer32 (0..65535)
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "This object is the configuration of the TPID used
        in the Upstream Direction (US) by to translate
        the TPID of the frame.
        Value of 0 indicates that no TPID translation is required in the
        Upstream Direction."
    ::= { docsL2vpnCmNsiEntry 8 }

-----
--
-- Point-to-Multipoint Only
--
-- The following objects are required for Point-to-Multipoint
-- operation only.
--
-----
--
-- Cable Modem/Vpn/CPE Table

```

```

-- (Point-to-Multipoint only)
--

docsL2vpnCmVpnCpeTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnCmVpnCpeEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "This table is a list of CPEs, indexed by the VPNs on a
        Cable Modem."
    ::= { docsL2vpnMIBObjects 10 }

docsL2vpnCmVpnCpeEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnCmVpnCpeEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "This table is a list of CPEs, indexed by the VPNs on a
        Cable Modem."
    INDEX { docsIfCmtsCmStatusIndex,
            docsL2vpnIdx,
            docsL2vpnCmVpnCpeMacAddress }
    ::= { docsL2vpnCmVpnCpeTable 1 }

DocsL2vpnCmVpnCpeEntry ::= SEQUENCE {
    docsL2vpnCmVpnCpeMacAddress  MacAddress
}

docsL2vpnCmVpnCpeMacAddress OBJECT-TYPE
    SYNTAX      MacAddress
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The Customer Premise Equipment (CPE) Mac Address
        that is attached to this instance's Cable Modem
        and bridging on this instance's VPN Id."
    ::= { docsL2vpnCmVpnCpeEntry 1 }

-----
--
-- VPN/Cable Modem/CPE Table
-- (Point-to-Multipoint only)
--
docsL2vpnVpnCmCpeTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnVpnCmCpeEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "This table contains a list of bridging CPEs, indexed by
        L2VPN Index and the corresponding CMs on that VPN."
    ::= { docsL2vpnMIBObjects 11 }

docsL2vpnVpnCmCpeEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnVpnCmCpeEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "This table contains a list of bridging CPEs, indexed by
        VPN and the corresponding CMs on that VPN."
    INDEX { docsL2vpnIdx,
            docsIfCmtsCmStatusIndex,
            docsL2vpnVpnCmCpeMacAddress }
    ::= { docsL2vpnVpnCmCpeTable 1 }

```

```

DocsL2vpnVpnCmCpeEntry ::= SEQUENCE {
    docsL2vpnVpnCmCpeMacAddress    MacAddress
}

docsL2vpnVpnCmCpeMacAddress OBJECT-TYPE
    SYNTAX      MacAddress
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The Customer Premise Equipment (CPE) Mac Address
         that is attached to this instance's Cable Modem
         and bridging on this instance's L2vpn Index."
    ::= { docsL2vpnVpnCmCpeEntry 1 }

-----
--
-- dot1qTpFdbTable Extension
-- (Point-to-Multipoint only)
--
docsL2vpnDot1qTpFdbExtTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnDot1qTpFdbExtEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This table contains packet counters for
         Unicast MAC Addresses within a VPN."
    ::= { docsL2vpnMIBObjects 12 }

docsL2vpnDot1qTpFdbExtEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnDot1qTpFdbExtEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This table extends the dot1qTpFdbTable only for RF network
         bridge port entries. It is implemented by an agent only
         if the agent implements dot1qTpFdbTable for RF network
         bridge ports."
    INDEX { dot1qFdbId, dot1qTpFdbAddress }
    ::= { docsL2vpnDot1qTpFdbExtTable 1 }

DocsL2vpnDot1qTpFdbExtEntry ::= SEQUENCE {
    docsL2vpnDot1qTpFdbExtTransmitPkts    Counter32,
    docsL2vpnDot1qTpFdbExtReceivePkts     Counter32
}

docsL2vpnDot1qTpFdbExtTransmitPkts OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The number of packets where the Destination
         MAC Address matched this instance
         dot1qTpFdbAddress and packet was bridged on
         a VPN, where the VPN ID matched this
         instance's dot1qFdbId."
    ::= { docsL2vpnDot1qTpFdbExtEntry 1 }

docsL2vpnDot1qTpFdbExtReceivePkts OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION

```

```

        "The number of packets where the Source MAC
        Address matched this instance dot1qTpFdbAddress
        and the packet was bridged on a VPN,
        where the docsL2vpnIdx matched this instance's
        dot1qFdbId."
    ::= { docsL2vpnDot1qTpFdbExtEntry 2 }

-----
--
-- dot1qTpGroupTable Extension
-- (Point-to-multipoint only)
--
docsL2vpnDot1qTpGroupExtTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnDot1qTpGroupExtEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This table contains packet counters for
        Multicast MAC Addresses within a VPN."
    ::= { docsL2vpnMIBObjects 13 }

docsL2vpnDot1qTpGroupExtEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnDot1qTpGroupExtEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This table extends the dot1qTpGroupTable only for RF
        Network bridge port entries. It is implemented by an agent
        Only if the agent implements dot1qTpGroupTable for RF
        network bridge ports."
    INDEX { dot1qVlanIndex, dot1qTpGroupAddress }
    ::= { docsL2vpnDot1qTpGroupExtTable 1 }

DocsL2vpnDot1qTpGroupExtEntry ::= SEQUENCE {
    docsL2vpnDot1qTpGroupExtTransmitPkts    Counter32,
    docsL2vpnDot1qTpGroupExtReceivePkts     Counter32
}

docsL2vpnDot1qTpGroupExtTransmitPkts OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The number of packets where the Destination
        MAC Address matched this instance
        dot1qTpGroupAddress and packet was bridged on
        a VPN, where the docsL2vpnIdx matched this
        instance's dot1qVlanIndex."
    ::= { docsL2vpnDot1qTpGroupExtEntry 1 }

docsL2vpnDot1qTpGroupExtReceivePkts OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The number of packets where the Source MAC
        Address matched this instance dot1qTpGroupAddress
        and the packet was bridged on a VPN,
        where the docsL2vpnIdx matched this instance's
        dot1qVlanIndex."
    ::= { docsL2vpnDot1qTpGroupExtEntry 2 }

```

```

-----
--
-- L2VPN/SF Statistics Table
-- Point-to-Point and Multipoint mode
--
docsL2vpnVpnSfStatsTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF DocsL2vpnVpnSfStatsEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "This table contains statistics for forwarding of
        packets on each VPN Service Flow."
    ::= { docsL2vpnMIBObjects 14 }

docsL2vpnVpnSfStatsEntry OBJECT-TYPE
    SYNTAX      DocsL2vpnVpnSfStatsEntry
    MAX-ACCESS  not-accessible
    STATUS      current
    DESCRIPTION
        "An entry is indexed by ifIndex of the CMTS MAC interface
        and Service Flow Index. The ifIndex is of type
        docsCableMacLayer(127)"
    INDEX { docsL2vpnIdx, docsQosServiceFlowId }
    ::= { docsL2vpnVpnSfStatsTable 1 }

DocsL2vpnVpnSfStatsEntry ::= SEQUENCE {
    docsL2vpnVpnSfStatsUpstreamPkts      Counter32,
    docsL2vpnVpnSfStatsUpstreamBytes    Counter32,
    docsL2vpnVpnSfStatsUpstreamDiscards Counter32,
    docsL2vpnVpnSfStatsDownstreamPkts   Counter32,
    docsL2vpnVpnSfStatsDownstreamBytes  Counter32,
    docsL2vpnVpnSfStatsDownstreamDiscards Counter32
}

docsL2vpnVpnSfStatsUpstreamPkts OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The number of L2vpn-forwarded packets received
        from this instance's Service Flow on
        this instance's L2VPN."
    ::= { docsL2vpnVpnSfStatsEntry 1 }

docsL2vpnVpnSfStatsUpstreamBytes OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The number of L2vpn-forwarded bytes received
        from this instance's Service Flow on
        this instance's L2VPN."
    ::= { docsL2vpnVpnSfStatsEntry 2 }

docsL2vpnVpnSfStatsUpstreamDiscards OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The number of L2-forwarded packets
        discarded from this instance's
        Service Flow on this instance's VPN."
    ::= { docsL2vpnVpnSfStatsEntry 3 }

```

```

docsL2vpnVpnSfStatsDownstreamPkts OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The number of L2-forwarded packets
        transmitted to this instance's
        Service Flow on this instance's VPN."
    ::= { docsL2vpnVpnSfStatsEntry 4 }

docsL2vpnVpnSfStatsDownstreamBytes OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The number of L2-forwarded bytes
        transmitted to this instance's
        Service Flow on this instance's VPN."
    ::= { docsL2vpnVpnSfStatsEntry 5 }

docsL2vpnVpnSfStatsDownstreamDiscards OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The number of L2-forwarded packets that were discarded
        before they could be transmitted to this instance's
        Service Flow on this instance's VPN."
    ::= { docsL2vpnVpnSfStatsEntry 6 }

-----

--
-- Conformance definitions
--
docsL2vpnConformance OBJECT IDENTIFIER ::= { docsL2vpnMIB 2 }
docsL2vpnCompliances OBJECT IDENTIFIER ::= { docsL2vpnConformance 1 }
docsL2vpnGroups OBJECT IDENTIFIER ::= { docsL2vpnConformance 2 }

docsL2vpnCompliance MODULE-COMPLIANCE
    STATUS       current
    DESCRIPTION
        "The compliance statement for the Cable Modem Termination
        Systems that implement the DOCSIS L2VPN Feature."

    MODULE      -- docsL2vpn
        -- conditionally mandatory groups
    GROUP docsL2vpnBaseGroup
    DESCRIPTION
        "Mandatory in all CMTSS."

    GROUP docsL2vpnPointToPointGroup
    DESCRIPTION
        "Mandatory in all CMTSS that implement point-to-point L2VPN
        forwarding."

    GROUP docsL2vpnMultipointGroup
    DESCRIPTION
        "Mandatory in all CMTSS that implement Multipoint
        L2VPN Forwarding Mode for any L2VPN."

    ::= { docsL2vpnCompliances 1 }

```

```

docsL2vpnBaseGroup OBJECT-GROUP
  OBJECTS {
    docsL2vpnIdToIndexIdx,
    docsL2vpnIndexToIdId,

    docsL2vpnCmCompliantCapability,
    docsL2vpnCmDutFilteringCapability,
    docsL2vpnCmDutCMIM,
    docsL2vpnCmDhcpSnooping,

    docsL2vpnVpnCmCMIM,
    docsL2vpnVpnCmVendorSpecific,
    docsL2vpnVpnCmIndividualSAId,

    docsL2vpnVpnCmStatsUpstreamPkts,
    docsL2vpnVpnCmStatsUpstreamBytes,
    docsL2vpnVpnCmStatsUpstreamDiscards,
    docsL2vpnVpnCmStatsDownstreamPkts,
    docsL2vpnVpnCmStatsDownstreamBytes,
    docsL2vpnVpnCmStatsDownstreamDiscards,

    docsL2vpnPortStatusGroupSAId,

    docsL2vpnSfStatusL2vpnId,
    docsL2vpnSfStatusUpstreamUserPriority,
    docsL2vpnSfStatusVendorSpecific,

    docsL2vpnPktClassL2vpnIdx,
    docsL2vpnPktClassUserPriRangeLow,
    docsL2vpnPktClassUserPriRangeHigh,
    docsL2vpnPktClassCMIM,
    docsL2vpnPktClassVendorSpecific,

    docsL2vpnVpnSfStatsUpstreamPkts,
    docsL2vpnVpnSfStatsUpstreamBytes,
    docsL2vpnVpnSfStatsUpstreamDiscards,
    docsL2vpnVpnSfStatsDownstreamPkts,
    docsL2vpnVpnSfStatsDownstreamBytes,
    docsL2vpnVpnSfStatsDownstreamDiscards
  }
  STATUS      current
  DESCRIPTION
    "A collection of objects in common for both
    Point-to-Point and Multipoint L2VPN forwarding
    Modes."
  ::= { docsL2vpnGroups 1 }

docsL2vpnPointToPointGroup OBJECT-GROUP
  OBJECTS {
    docsL2vpnCmNsiEncapSubtype,
    docsL2vpnCmNsiEncapValue,
    docsL2vpnCmNsiAGI,
    docsL2vpnCmNsiSAII,
    docsL2vpnCmNsiTAII,
    docsL2vpnCmNsiMode,
    docsL2vpnCmNsiTpidTransDS,
    docsL2vpnCmNsiTpidTransUS
  }
  STATUS      current
  DESCRIPTION

```

```
        "A collection of objects in common for only the
        Point-to-Point forwarding mode."
    ::= { docsL2vpnGroups 2 }

docsL2vpnMultipointGroup OBJECT-GROUP
    OBJECTS {
        docsL2vpnCmVpnCpeMacAddress,

        docsL2vpnVpnCmCpeMacAddress,

        docsL2vpnDot1qTpFdbExtTransmitPkts,
        docsL2vpnDot1qTpFdbExtReceivePkts,

        docsL2vpnDot1qTpGroupExtTransmitPkts,
        docsL2vpnDot1qTpGroupExtReceivePkts
    }
    STATUS      current
    DESCRIPTION
        "A collection of objects required only for Multipoint
        forwarding mode."
    ::= { docsL2vpnGroups 3 }

END
```

Annex B Parameter Encodings (Normative)

B.1 Capabilities

B.1.1 L2VPN Capability

This capability indicates whether the CM is compliant with the Layer 2 Virtual Private Network requirements for a CM that are specified in Section 7. L2VPN operation may still be performed with CMs that do not implement these requirements, but with possible limitations.

Type	Length	Value
5.17	1	0 CM not compliant with DOCSIS L2VPN Section 7 (default) 1 CM compliant with DOCSIS L2VPN Section 7

B.1.2 Embedded Service/Application Functional Entity (eSAFE) Host Capability ⁶⁵

This capability encoding informs the CMTS of the type and MAC address of an eSAFE host embedded with the CM. This is necessary for the CMTS to guarantee proper IPv4 and IPv6 forwarding of upstream traffic from the eSAFE host. A separate eSAFE Host Capability encoding is required for each separate eSAFE host embedded with the CM.

Type	Length	Value
5.18	7	eSAFE ifIndex (1 byte), eSAFE MAC address (6 bytes) eSAFE ifIndex: 1 eRouter 16 eMTA 17 eSTB-IP 19 eTEA

B.1.3 Downstream Unencrypted Traffic (DUT) Filtering

This capability indicates whether the CM supports the DUT Filtering feature as described in Section 6.5.2.1.

Type	Length	Value
5.19	1	0 DUT Filtering not supported (default) 1 DUT Filtering supported

B.2 Downstream Unencrypted Traffic (DUT) Filtering Encoding

The DUT Filtering parameter is intended for CMs implementing layer 2 or layer 3 Virtual Private Networks. In such networks, downstream traffic intended for the private network is always encrypted with BPI. Because the RF downstream is broadcast to all CMs, however, unencrypted group MAC traffic intended for *other* CMs will leak onto the VPN CM's CMCI ports unless filtered in the VPN CM. This parameter allows VPN CMs to filter all downstream unencrypted traffic, to both individual and group MAC destinations.

⁶⁵ Section modified per L2VPN-N-07.0504-1 and L2VPN-N-07.0538-2 on 11/9/07 by KN; TLV Table modified per L2VPN-N-12.1065-11 on 9/19/12 by CP and L2VPN-N-13.1120-2 on 11/7/13 by PO.

Top-Level TLV Encoding for DUT Filtering:⁶⁶

Type	Length	Value
45	N	Composite value ⁶⁷

The encapsulated fields are described below.

B.2.1 Downstream Unencrypted Traffic (DUT) Control

Type	Length	Value
45.1	1	Bit 0 DUT Filtering DUT Filtering = 0: Disable (default) DUT Filtering = 1: Enable DUT Filtering. Bits 1..7: Reserved. ⁶⁸

If the DUT Filtering Encoding is omitted, or the DUT filtering bit is zero, then the CM bridges downstream unencrypted traffic, received from its RF interface (ifIndex 2) according to relevant DOCSIS specifications, namely forwarding unicast MAC traffic to the internal (eCM/eSAFE) or external (CPE) bridge port from which a source MAC was learned or configured and forwarding IGMP-learned or configured group MAC (GMAC) traffic to all other internal and external bridge ports.⁶⁹

If the DUT Filtering Encoding is present and the DUT Filtering bit is set to "1", the CM MUST restrict forwarding of downstream unencrypted traffic (for both individual and group MAC destinations) to only the set of interfaces indicated in a DUT CM Interface Mask (DUT CMIM) configured or implied by the encoding.

B.2.2 Downstream Unencrypted Traffic (DUT) CMIM

Type	Length	Value
45.2	N	DUT CMIM, optional CM Interface Mask (CMIM) limiting outgoing interfaces of DUT traffic. If the DUT CMIM is omitted, its default value includes the eCM and all implemented eSAFE interfaces, but not any CPE interfaces.

An explicit DUT CMIM subtype has the format as defined in Section B.3.4.

If this subtype is not present the implied DUT CMIM includes the eCM (CMIM bit position 0) and all implemented eSAFE interfaces on the CM (CMIM bit positions 16 and higher), but excludes all CPE interfaces.⁷⁰ This implied DUT CMIM permits a cable operator to configure a DUT Filtering Encoding that is generic for all CM device types offering Transparent LAN Service.

B.3 L2VPN Encoding

The L2VPN Encoding parameter is a multi-part encoding that configures how the CMTS performs Layer 2 Virtual Private Network bridging for CPE packets. L2VPN operation is specified in Section 6.

An L2VPN Encoding is termed a per-SF L2VPN Encoding when it appears as a subtype of the Upstream Service Flow Encoding (type 24). The encoding is a Downstream Classifier L2VPN Encoding when it appears in a Downstream Packet Classification Configuration Setting (type 23). It is termed an Upstream Classifier L2VPN encoding when it appears in an Upstream Packet Classification Configuration Setting (type 22).

⁶⁶ Modified per L2VPN-N-07.0407-1 on 4/20/07 by KN; modified by L2VPN-N-14.1186-1 on 1/15/15 by PO.

⁶⁷ Modified per L2VPN-N-06.0277-1 on 7/26/06 by KN.

⁶⁸ Previous sentence and table added per L2VPN-N-06.0277-1 on 7/26/06 by KN.

⁶⁹ Modified per L2VPN-N-06.0277-1 on 7/26/06 by KN.

⁷⁰ Text modified and previous table added per L2VPN-N-06.0277-1 on 7/26/06 by KN.

A Forwarding L2VPN Encoding is one that contains an L2VPN VPNID subtype that configures forwarding of packets on a particular L2VPN.

The L2VPN Encoding is encoded as General Extension Information (GEI) ([RFI 2.0]) as a Vendor Specific encoding with vendor ID 0xFFFFF. GEI subtype 5 is assigned for L2VPN Encodings.

GEI-Encapsulated L2VPN Encoding:

Type	Length	Value
43.5	n	L2VPN subtype/length/value tuples

The L2VPN Encoding itself contains one or more L2VPN subtype encodings.

B.3.1 VPN Identifier ⁷¹

The VPN Identifier (VPNID) subtype encoding is an opaque octet string identifier that associates an attachment circuit (i.e., a CM or one SF of a CM) or a Downstream Classifier to a particular layer 2 virtual private network. VPN ID values are advantageously configured as printable ASCII strings. VPNID values are unique within a single CMTS. VPNID values are advantageously configured as unique across *all* CMTSs within the administrative domain of the cable operator operating the CMTS. VPNID values can be configured to be *globally* unique in order to facilitate inter-domain L2VPN forwarding.

For scalability, a cable operator can configure VPNID strings to algorithmically map to globally-unique binary octet strings. Examples of globally-unique binary octet strings for VPNs include the 7-byte VPN ID format described in [RFC 2685] and the 8-byte Route Distinguisher described in [RFC 4364]. A single VPNID subtype is present in valid Forwarding L2VPN Encodings.

In general, multiple attachment circuits (i.e., CM/SFs) can connect to the same L2VPN, and so would be configured with the same VPNID subtype value. If the CMTS performs only point-to-point L2VPN forwarding for the indicated L2VPN without using any auto-discovery protocol, it enforces that the L2VPN Encoding also contains an NSI Encapsulation Subtype.

A CMTS performing Multipoint L2VPN Forwarding MUST perform transparent learning layer 2 forwarding between 802.1Q bridge ports, cable modems, and service flows configured with the same VPNID.

In per-SF L2VPN Encodings, the VPNID identifies the L2VPN on which upstream traffic is to be forwarded. In Downstream Classifier L2VPN Encodings within a Downstream Packet Classification Setting, the VPNID configures the classifier to apply to only L2VPN-forwarded downstream traffic on the L2VPN identified by the VPNID.

A CMTS SHOULD use the value of the VPNID with any signaling protocols that dynamically determine Service Multiplexing field values on L2VPN packets encapsulated on an NSI port. The VPNID is intended to be (or map to) the attachment group identifier (AGI) for IETF L2VPN working group signaling protocols.

The CMTS MUST support configuration of VPNID values of at least 4 octets, and no more than 255 octets. The number of unique VPNID values supported by the CMTS is vendor-specific.

Note: The VPNID is NOT used to identify an L2VPN for BGP auto-discovery ([RFC 6074]). BGP Auto-Discovery uses the BGP VPNID (43.5.21.1, see Section B.3.21.1).

Sub-Type	Length	Value
43.5.1	1..N	An opaque octet string that identifies a Layer 2 Virtual Private Network. N is vendor specific, but has to be within the range 4..255.

⁷¹ Various changes throughout section per L2VPN-N-12.1065-11 on 9/19/12 by CP.

B.3.2 NSI Encapsulation Subtype⁷²

At a minimum, this subtype is required only to specify how the CMTS encapsulates Point-to-Point L2VPN-forwarded packets and forwards accordingly over the NSI port, primarily for L2VPN feature certification testing. It is intended, however, to also allow cable operator configuration of IETF Pseudo Wire Emulation [RFC 3985] of each cable attachment circuit (CM or SF) across the NSI backbone when BGP auto-discovery is not to be used by the CMTS.⁷³

In Selected Ethernet mode, the CMTS is configured to forward all L2VPN traffic through a single Ethernet NSI port at any given time. When a Selected Ethernet Port is identified, the CMTS MUST accept the IEEE 802.1Q NSI Encapsulation Format Code in Forwarding L2VPN Encodings and MAY accept the other codes.

Although the NSI Encapsulation Subtype is intended primarily for Point-to-Point forwarding modes, the CMTS may accept it in Multipoint mode (including in the Selected Ethernet mode). If the NSI Encapsulation Subtype is specified, the CMTS MUST accept and implement it, provided that it does not differ from any NSI Encapsulation Subtype for that VPNID within any other accepted L2VPN Encoding.⁷⁴

The value of the NSI Encapsulation Subtype is a single Format Code-Length-Value tuple that identifies an NSI Encapsulation Format Code and possibly, an NSI Encapsulation Service Multiplexing value.

Sub-Type	Length	Value
43.5.2	n	A single NSI encapsulation format code/length/value tuple

If the NSI Encapsulation Subtype or an L2VPN Vendor Specific Subtype does not statically configure a Service Multiplexing value, the CMTS MUST dynamically select and learn the Service Multiplexing value for a forwarding L2VPN Encoding from the CMTS's L2VPN peers across the NSI interface. Dynamically learned Service Multiplexing values may be different on different NSI ports.

NSI Encapsulation Format Code ⁷⁵	Length	Service Multiplexing Value
43.5.2.1	0	Other. The L2VPN NSI Encapsulation format is other than those specified below. In this case, L2VPN Vendor Specific Subtype Encodings (GEI Subtype 5.43) MUST provide the NSI Encapsulation Format and any desired static Service Multiplexing values.
43.5.2.2	2	IEEE 802.1Q. Value is the 16-bit IEEE 802.1Q tag (most significant byte first) that contains, in its least significant 12 bits, a VLAN ID used to recognize packets for the L2VPN on the Selected Ethernet NSI port. The most significant 4 bits of the 16-bit tag value are reserved. The CMTS SHOULD ignore the most significant 4 bits of the 16-bit NSI Encapsulation IEEE 802.1Q tag value. The maximum number of unique VLAN ID values accepted by a CMTS is vendor specific. A CMTS MUST accept the full 12-bit range of VLAN ID values for the unique values it does accept.
43.5.2.3	4	IEEE 802.1Q (formerly known as 802.1 ad). Value is a pair of 16-bit values (most significant byte first), with the first 16-bit field containing a Service Provider VLAN ID in the least significant 12 bits, and the second 16-bit field containing the Customer VLAN ID in the least significant 12 bits. The most significant 4 bits of the first 16-bit value is used as follows: Three MSB bits: S-PCP, the remaining 1 bit: S-DEI The most significant 4 bits of the second 16-bit value is used as follows: Three MSB bits: C-PCP, the remaining 1 bit: C-CFI The maximum number of Service Provider and Customer VLAN ID values the CMTS accepts is vendor specific, but the CMTS MUST accept the full 12-bit range of VLAN ID values.

⁷² Revised per L2VPN-N-10.0918-2 on 5/10/10 by JB, L2VPN-N-13.1109-3 on 7/23/13 by PO, replacement of 802.1ah, 802.1ad, and 802.1ag throughout this section per L2VPN-N-14.1186-1 on 1/15/15 by PO.

⁷³ Paragraph revised per L2VPN-N-12.1065-11 on 9/19/12 by CP.

⁷⁴ Paragraph modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.

⁷⁵ Table revised per L2VPN-N-12.1065-11 on 9/19/12 by CP.

NSI Encapsulation Format Code ⁷⁵	Length	Service Multiplexing Value
43.5.2.4	n	MPLS PW. This TLV defines the primary MPLS PW Identifier and PW Peer IP address and optionally a Pseudowire Type, Backup MPLS PW ID and PW Peer IP address to be used by the CMTS. The attachment circuit's L2VPN traffic is intended to forward over an MPLS label switched path to the peer. The CMTS SHOULD dynamically select and learn the label stack for incoming and outgoing label stacks, respectively. The CMTS MAY use Vendor Specific L2VPN Subtypes to statically configure the ingress and egress label stacks. The CMTS MAY limit statically configured MPLS label values to a vendor-specific range.
43.5.2.4.1	4	MPLS Pseudowire ID. Value is a 4-byte Identifier in range of 1-4294967296
43.5.2.4.2	5 or 17	MPLS Peer IP address. Value is a 1-byte InetAddressTypeCode (ipv4(1) or ipv6(2)) followed by an IPv4 or IPv6 InetAddress 1st byte: 1: IPv4, 2: IPv6 peer bytes 2-5 or 2-17: peer IP address
43.5.2.4.3	1	Pseudowire Type. Value is either 4 (Ethernet Tagged Mode) or 5 (Ethernet Raw Mode) or 19 (VPLS) [PWYPES]. If absent, then value of 5 should be assumed by the CMTS. If present, the CMTS SHOULD use it as the Pseudowire Type in the PW signaling (e.g., LDP).
43.5.2.4.4	4	MPLS Backup Pseudowire ID. Value is a 4-byte Identifier in range of 1-4294967296 This TLV defines the Backup MPLS PW ID to be used by the CMTS to activate the Pseudowire Redundancy. When the 43.5.2.4.4 TLV is present, the CMTS MUST provision the Backup MPLS PW for the attachment circuit.
43.5.2.4.5	5 or 17	MPLS Backup Peer IP address. Value is a 1-byte InetAddressTypeCode (ipv4(1) or ipv6(2)) followed by an IPv4 or IPv6 InetAddress 1st byte: 1: IPv4, 2: IPv6 peer bytes 2-5 or 2-17: peer IP address This TLV defines the Backup MPLS PW Peer to be used by the CMTS to activate the Pseudowire Redundancy. When the 43.5.2.4.5 TLV is present, the CMTS MUST provision the Backup MPLS PW for the primary PW.
43.5.2.5	5 or 17	L2TPv3 Peer. Value is a one-byte InetAddressTypeCode (ipv4(1) or ipv6(2)) followed by an IPv4 or IPv6 InetAddress. The attachment circuit's L2VPN traffic is intended to forward within an L2TPv3 tunnel to the addressed peer. The CMTS SHOULD dynamically select and learn the local and remote session IDs for each tunnel. The CMTS MAY use Vendor Specific L2VPN Subtypes to statically configure session IDs, L2TPv3 peer network addresses, and other information as required by the vendor. The CMTS MAY limit statically configured session ID or other Service Multiplexing values to a vendor-specific range.
43.5.2.6	n	IEEE 802.1Q Encapsulation. This parameter defines the parameters associated with [802.1Q] tagging and encapsulation. Two subtypes are required.
43.5.2.6.1	4	IEEE 802.1Q Backbone Service Instance Tag (I-Tag) TCI. This parameter defines [802.1Q] I-Tag TCI field, which consists of TPID 0x88e7 and this 32-bit I-Tag TCI (most significant byte 1st). The TCI value contains in its least significant 24-bits the Backbone Service Instance Identifier (I-SID). The most significant byte of I-Tag TCI has to be zero. 32-bit value of [802.1Q] I-Tag TCI
43.5.2.6.2	6	IEEE 802.1Q Destination Backbone Edge Bridge (BEB) MAC Address (B-DA). This parameter defines for a given Backbone Service Instance the MAC address of the destination BEB, which should deliver [802.1Q] frames of this instance to the destination customer systems. The value of this parameter is 6-bytes individual MAC address. 48-bit BEB MAC Address
43.5.2.6.3	2	16-bit value of [802.1Q] B-Tag TCI
43.5.2.6.4	2	16-bit value of [802.1Q] I-Tag TPID
43.5.2.6.5	1	3 bit I-PCP
43.5.2.6.6	1	1 bit I-DEI
43.5.2.6.7	1	1 bit I-UCA
43.5.2.6.8	3	24-bit value of [802.1Q] I-SID Backbone Service Instance Identifier
43.5.2.6.9	2	16-bit value of [802.1Q] B-Tag TPID
43.5.2.6.10	1	1 bit B-PCP

NSI Encapsulation Format Code ⁷⁵	Length	Service Multiplexing Value
43.5.2.6.11	1	1 bit B-DEI
43.5.2.6.12	2	12-bit value of [802.1Q] B-VID
43.5.2.8	2	16-bit value of [802.1Q] S-TPID

B.3.2.1 [802.1Q] S-TPID (formerly known as 802.1ad)⁷⁶

This TLV defines [802.1Q] S-TPID value is to be used in the [802.1Q] Encapsulation Mode. If this TLV is not specified, a default value 0x88a8 will be used for [802.1Q] S-TPID field.

SubType	Length	Value
43.5.2.8	2	16-bit value of [802.1Q] S-TPID

B.3.2.2 [802.1Q] Encapsulation (formerly known as 802.1ah)

This TLV defines the parameters associated with [802.1Q] encapsulation.

SubType	Length	Value
43.5.2.6	n	

B.3.2.2.1 [802.1Q] I-TCI

This TLV defines the value to be used for 32-bit [802.1Q] I-Tag TCI (most significant byte 1st), which contains 3 bits I-PCP, 1 bit I-DEI, 1 bit I-UCA, 3 bits Reserved and least significant 24-bits the Backbone Service Instance Identifier (I-SID).

SubType	Length	Value
43.5.2.6.1	4	32-bit value of [802.1Q] I-Tag TCI

B.3.2.2.2 MAC Address of the Destination Backbone Edge Bridge (B-DA)

This TLV defines for a given Backbone Service Instance the MAC address of the destination BEB, which should deliver [802.1Q] frames of this instance to the destination customer systems. The value of this TLV is 6-bytes individual MAC address.

SubType	Length	Value
43.5.2.6.2	6	48-bit BEB MAC Address

B.3.2.2.3 [802.1Q] B-TCI

[802.1Q] B-Tag consists of TPID 0x88a8 and 16-bit B-TAG TCI. This TLV defines the value of the 16-bit [802.1Q] B-Tag TCI (most significant byte first), which contains 3 bits B-PCP, 1-bit B-DEI and least significant 12-bits the Backbone Service Instance Identifier (B-VID).

SubType	Length	Value
43.5.2.6.3	2	16-bit value of [802.1Q] B-Tag TCI

⁷⁶ This subsection and the following subsections up to B.3.3 added per L2VPN-N-12.1065-11 on 9/19/12 by CP.

B.3.2.2.4 [802.1Q] I-TPID

This TLV defines [802.1Q] I-TPID value to be used in the [802.1Q] Encapsulation Mode. If this TLV is not specified, a default value 0x88e7 will be used for [802.1Q] I-TPID field.

SubType	Length	Value
43.5.2.6.4	2	16-bit value of [802.1Q] I-Tag TPID

B.3.2.2.5 [802.1Q] I-PCP

This TLV defines [802.1Q] I-PCP value to be used in the [802.1Q] Encapsulation Mode. If this TLV is not specified, a default value 0 will be used for [802.1Q] I-PCP field.

SubType	Length	Value
43.5.2.6.5	1	This TLV comprises an encoded bit map, featuring one field: I-PCP, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	5 bits
I-PCP	Encodes the I-Tag PCP field	3 bits

B.3.2.2.6 [802.1Q] I-DEI

This TLV defines [802.1Q] I-DEI value to be used in the [802.1Q] Encapsulation Mode. If this TLV is not specified, a default value 0 will be used for [802.1Q] I-DEI field.

SubType	Length	Value
43.5.2.6.6	1	This TLV comprises an encoded bit map, featuring one field: I-DEI, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	7 bits
I-DEI	Encodes the I-DEI field	1 bit

B.3.2.2.7 [802.1Q] I-UCA

This TLV defines [802.1Q] I-UCA value to be used in the [802.1Q] Encapsulation Mode. If this TLV is not specified, a default value 0 will be used for [802.1Q] I-UCA field.

SubType	Length	Value
43.5.2.6.7	1	This TLV comprises an encoded bit map, featuring one field: I-UCA, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	7 bits
I-UCA	Encodes the I-UCA field	1 bit

B.3.2.2.8 [802.1Q] I-SID

This TLV defines 24-bits [802.1Q] I-SID Backbone Service Instance Identifier value to be used in the [802.1Q] Encapsulation Mode.

SubType	Length	Value
43.5.2.6.8	3	24-bit value of [802.1Q] I-SID Backbone Service Instance Identifier

B.3.2.2.9 [802.1Q] B-TPID

This TLV defines [802.1Q] B-TPID value to be used in the [802.1Q] Encapsulation Mode. If this TLV is not specified, a default value 0x88a8 will be used for [802.1Q] B-TPID field.

SubType	Length	Value
43.5.2.6.9	2	16-bit value of [802.1Q] B-Tag TPID

B.3.2.2.10 [802.1Q] B-PCP

This TLV defines [802.1Q] B-PCP value to be used in the [802.1Q] Encapsulation Mode. If this TLV is not specified, a default value 0 will be used for [802.1Q] B-PCP field.

SubType	Length	Value
43.5.2.6.10	1	This TLV comprises an encoded bit map, featuring one field: B-PCP, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	5 bits
B-PCP	Encodes the B-PCP field	3 bits

B.3.2.2.11 [802.1Q] B-DEI

This TLV defines [802.1Q] B-DEI value to be used in the [802.1Q] Encapsulation Mode. If this TLV is not specified, a default value 0 will be used for [802.1Q] B-DEI field.

SubType	Length	Value
43.5.2.6.11	1	This TLV comprises an encoded bit map, featuring one field: B-DEI, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	7 bits
B-DEI	Encodes the B-DEI field	1 bit

B.3.2.2.12 [802.1Q] B-VID

This TLV defines [802.1Q] B-VID value to be used in the [802.1Q] Encapsulation Mode.

SubType	Length	Value
43.5.2.6.12	2	This TLV comprises an encoded bit map, featuring one field: B-VID, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	4 bits
B-VID	Encodes the B-VID field	12 bits

B.3.3 eSAFE DHCP Snooping

This parameter is defined only in a per-SF forwarding L2VPN encoding. The parameter is a bit mask with bit positions defined for each potential eSAFE host type. A '1' in the eSAFE host type's bit position enables the CMTS to automatically detect the MAC address of that eSAFE host by snooping DHCP traffic forwarded between the CM and a DHCP server. The bit positions in the eSAFE DHCP Snooping parameter match those of the CM Interface Mask (CMIM) for the interface associated with the eSAFE host type.

SubType	Length	Value
43.5.3	1..N	Bit mask of eSAFE hosts enabled for DHCP Snooping Bit 1 (0x40 00 00): eRouter ⁷⁷ Bit 16 (0x00 00 80) PacketCable-EMTA Bit 17 (0x00 00 40) eSTB-IP Bit 18 (0x00 00 20) eSTB-DSG Bits 19..31 (0x00 00 1F FF) Other eSAFE interfaces

B.3.4 CM Interface Mask (CMIM) Subtype

This parameter is a bit mask that describes a set of eCM interface indexes [eDOCSIS]. In a Forwarding L2VPN Encoding, the CM Interface Mask Subtype describes the set of bridge port interfaces on which the CM forwards packets of the L2VPN.

Each bit of CMIM corresponds to a logical bridge port interface of a MAC layer 2 bridge implemented in the eCM of a cable modem. The parameter is encoded as the octet string of the Basic Encoding Rules encoding of an SNMP BITS bit string. Bit position K in the BITS encoding corresponds to eDOCSIS MAC bridge interface K. By convention, bit position 0 corresponds to the eCM's self host interface. The eCM self MAC address is signaled as if it were on a bridge port interface ifIndex of zero (0), even though no such interface actually exists.

SubType	Length	Value
43.5.4	N	SNMP BITS -encoded bit map with bit position K representing eCM logical interface index value K. Bit position 0 represents the eCM self host itself. Bit position 0 is the most significant bit of the first octet. Refer to [eDOCSIS] for latest logical interface index assignments. Bit 0 (0x80): eCM self host interface Bit 1 (0x40): primary CPE Interface (also eRouter) ⁷⁸ Bit 2 (0x20) RF interface Bits 3,4 reserved Bits 5..15 (0x07 FF) Other CPE Interfaces Bits 16-31, embedded logical interfaces. Currently defined interfaces include: Bit 16 (0x00 00 80) PacketCable-EMTA Bit 17 (0x00 00 40) eSTB-IP Bit 18 (0x00 00 20) eSTB-DSG Bits 19..31 (0x00 00 1F FF) Other eSAFE interfaces

If the CM Interface Mask subtype is not present in a Forwarding L2VPN Encoding, its default value is for the Primary CPE Interface (index 1) and the cable RF interface (index 2) only, i.e., CMIM value 0x60. A CM MUST silently ignore CMIM bit positions for unimplemented interfaces. A CMTS MAY signal that a CMIM value represents all possible CPE interfaces with the CMIM value for positions 1 and 5-15, i.e., the CMIM value 0x47 FF.

⁷⁷ Revised per L2VPN-N-12.1065-11 on 9/19/12 by CP.

⁷⁸ per L2VPN-N-12.1065-11 on 9/19/12 by CP.

B.3.5 Attachment Group ID ⁷⁹

If present, the CMTS SHOULD use this subtype value as the Attachment Group ID (AGI) signaling element, associated with a VPN Identifier, when dynamically establishing an NSI pseudowire for Point-to-Point forwarding of the attachment circuit. It is applicable only in conjunction with MPLS or L2TPv3 NSI Encapsulation and Point-to-Point forwarding between the attachment circuit and the pseudowire.

AGI ID subtype should not be included if CMTS intends to use dynamic discovery method for discovering remote L2VPN forwarders (i.e., PE routers).

SubType	Length	Value
43.5.5	0..16	Byte string (opaque to the CM) having the following format: <AGI-type><AGI-value>, where AGI-type, encodes the AGI type as defined by IANA. It is of 1-byte length. AGI-value encodes the AGI value as defined by IANA for the preceding AGI-type. It is of variable length. For example, if AGI-type is 0x01, then AGI-value encodes 8-byte Route Distinguisher, resulting in a subtype TLV of length 9.

B.3.6 Source Attachment Individual ID ⁸⁰

If present, the CMTS SHOULD use this subtype value as the source attachment individual identifier (SAII) signaling element associated with the local pseudowire attachment when establishing an NSI backbone pseudowire for the cable attachment circuit. It is applicable only in conjunction with an MPLS or L2TPv3 NSI Encapsulation subtype and Point-to-Point forwarding between the cable attachment circuit and the pseudowire.

SAII ID subtype should not be included if CMTS intends to use dynamic discovery method for discovering remote L2VPN forwarders (i.e., PE routers).

SubType	Length	Value
43.5.6	0..16	Byte string (opaque to the CM) having the following format: <All-type><All-value>, where All-type encodes the All type as defined by IANA. It is of 1-byte length. All-value encodes the All value as defined by IANA for the preceding All-type. It is of variable length. For example, if All-type is 0x01, then All-value encodes 4-byte unsigned number, resulting in a subtype TLV of length 5.

B.3.7 Target Attachment Individual ID ⁸¹

If present, the CMTS SHOULD use this subtype value as the target attachment individual identifier (TAII) signaling element associated with the remote pseudowire attachment when establishing an NSI PseudoWire for the attachment circuit. It is applicable only in conjunction with an MPLS or L2TPv3 NSI Encapsulation subtype and Point-to-Point forwarding between the cable attachment circuit and the pseudowire.

TAII ID subtype should not be included if CMTS intends to use dynamic discovery method for discovering remote L2VPN forwarders (i.e., PE routers).

⁷⁹ Revised section per L2VPN-N-12.1065-11 on 9/19/12 by CP.

⁸⁰ Revised section per L2VPN-N-12.1065-11 on 9/19/12 by CP.

⁸¹ Revised section per L2VPN-N-12.1065-11 on 9/19/12 by CP.

SubType	Length	Value
43.5.7	0..16	<p>Byte string (opaque to the CM) having the following format: <All-type><All-value>., where</p> <p>All-type encodes the All type as defined by IANA. It is of 1-byte length.</p> <p>All-value encodes the All value as defined by IANA for the preceding All-type. It is of variable length.</p> <p>For example, if All-type is 0x01, then All-value encodes 4-byte unsigned number, resulting in a subtype TLV of length 5.</p>

B.3.8 Upstream User Priority ⁸²

IEEE 802.1D bridging protocols require the detection or generation, optional regeneration, and signaling of a user priority attribute of all bridged packets. The Upstream User Priority subtype is used to configure the [802.1D] user priority of outgoing L2VPN packets if the NSI encapsulation subtype is specified as 802.1Q. Similarly, the Upstream User Priority subtype is used to configure the MPLS Traffic Class (previously known as EXP) of outgoing L2VPN packets if the NSI encapsulation subtype is specified as MPLS(Pseudowire). This subtype is defined in only upstream per-SF Forwarding L2VPN Encodings.

Unless the L2VPN forwarder is otherwise configured, CMTS MUST transmit the upstream user priority signaled with this subtype either as the user priority bits of an IEEE 802.1Q tag when it forwards the packet to an NSI port with IEEE 802.1Q encapsulation, or as the Traffic Class bits (formerly known as EXP bits) of an MPLS label stack entry when it forwards the packet to an NSI port with MPLS(Pseudowire) encapsulation, or as the IP precedence bits of an IP packet when it forwards the packet to an NSI port with L2TPv3/IP encapsulation. If this subtype is omitted in a Forwarding L2VPN Encoding, the CMTS considers the incoming user priority to be zero (0). This subtype appears no more than once in a valid Forwarding L2VPN Encoding.

SubType	Length	Value
43.5.8	1	Ingress user priority value in the range 0..7 encoded in the least significant three bits of an [802.1Q] tag, the traffic class/EXP bits of an MPLS label stack, or the IP Precedence bits of an L2TPv3 packet header. Higher values indicate higher priority.

B.3.9 Downstream User Priority Range ⁸³

In a Downstream Packet Classification Encoding, the presence of an L2VPN Encoding with this subtype restricts the classifier to only packets forwarded downstream with the indicated range of user priority values (inclusive). The classified user priority is as transmitted on the DOCSIS MAC layer interface, and so is considered to be *after* any ingress default user priority selection or user priority regeneration performed by the L2VPN Forwarder. This subtype may appear only in a Downstream Classifier L2VPN Encoding, and at most, once in single L2VPN Encoding. If this subtype is omitted, the classifier applies to all egress user priority values.

SubType	Length	Value
43.5.9	2	Pri-low, pri-high. The lower user priority value of the user priority range is encoded in the least significant three bits of the first byte, and the higher value of the range is encoded in the least significant three bits of the second byte.

B.3.10 L2VPN SA-Descriptor Subtype

The CMTS adds this subtype in downstream Registration Response and Dynamic Service messages that contain forwarding L2VPN Encodings to inform an L2VPN-compliant CM of the SAID value(s) under which the CMTS

⁸² Revised per L2VPN-N-10.0918-2 on 5/10/10 by JB and per L2VPN-N-12.1065-11 on 9/19/12 by CP, Section modified per L2VPN-N-14.1186-1 on 1/15/15 by PO.

⁸³ Revised heading title per L2VPN-N-12.1065-11 on 9/19/12 by CP.

will encrypt the downstream traffic forwarded to that L2VPN through the CM. A valid L2VPN Encoding may have multiple L2VPN SA-Descriptor Subtypes.

SubType	Length	Value
43.5.10	14	SA-Descriptor compound subtypes specified in [BPI+] that provide the SAID value under which the CMTS encrypts downstream traffic forwarded an L2VPN. The SA-Type of the SA-Descriptor has to be Dynamic. ⁸⁴

B.3.11 Vendor Specific L2VPN Subtype

This subtype is interpreted by the CMTS in a vendor-specific fashion. An example usage is to configure the NSI sub-interface or virtual circuit to which upstream packets from the CM or SF are bridged in a Point-to-Point mode. The vendor-specific subtype contents can be binary or ASCII encoded data.

GEI Type	Length	Value
43.5.43	N	08, 3, vendor ID, following by vendor-specific type/length/value tuples.

B.3.12 Pseudowire ID (deprecated)⁸⁵

(This section is deprecated. See 43.5.2.4 for MPLS PW ID.)

B.3.13 Pseudowire Type⁸⁶

If present, the CMTS SHOULD use this subtype value as the Pseudowire Type associated with the local pseudowire attachment when establishing an NSI backbone pseudowire for the cable attachment circuit. It is applicable only in conjunction with an Ethernet over MPLS ([RFC 4448]) NSI Encapsulation subtype and Point-to-Point forwarding between the cable attachment circuit and the pseudowire. If implementing Ethernet over MPLS ([RFC 4448]), the CMTS MUST support Type 4 (Ethernet Tagged Mode) and Type 5 (Ethernet Raw Mode) pseudowires, as defined in [PWTYPES].

SubType	Length	Value
43.5.12	1	Pseudowire Type 4 (Ethernet Tagged Mode) or Type 5 (Ethernet Raw Mode) [PWTYPES].

B.3.14 L2VPN Mode⁸⁷

This parameter is used in DPoE specifications only. It is not used in DOCSIS L2VPN configuration.

This parameter is used to configure L2VPN Mode. L2VPN has two modes: Encapsulation Mode and Transport Mode. If this parameter is omitted or is 0, then Encapsulation Mode is used. In Encapsulation mode, the L2VPN NSI service multiplexing value, configured using the NSI Encapsulation Subtype (TLV 43.5.2), is to be used to add the service tag or encapsulation to frames associated with the given service.

If this parameter is set to "1", then Transport Mode is used. In Transport Mode, service tags are not added to or removed from frames associated with the given service.

⁸⁴ Table modified per L2VPN-N-06.0277 on 7/26/06 by KN and per L2VPN-N-12.1065-11 on 9/19/12 by CP.

⁸⁵ Revised per L2VPN-N-10.0918-2 on 5/10/10 by JB, then deprecated per L2VPN-N-12.1065-11 on 9/20/12 by CP.

⁸⁶ Revised per L2VPN-N-10.0918-2 on 5/10/10 by JB.

⁸⁷ Sections B.3.14 through B.3.26 added per L2VPN-N-12.1065-11 on 9/19/12 by CP.

SubType	Length	Value
43.5.13	1	L2VPN Mode 0 = L2VPN Encapsulation Mode (default) 1 = L2VPN Transport Mode Reserved

B.3.15 Tag Protocol Identifier (TPID) Translation

This parameter is used in DPoE specifications only. It is not used in DOCSIS L2VPN configuration.

This field defines the L2VPN top-level parameters associated with the translation of the TPID of the frame outermost tag.

SubType	Length	Value
43.5.14	n	

B.3.15.1 Upstream TPID Translation

This TLV defines the new TPID value to be used for the outermost Tag of the frame before it is transmitted on the NSI interface. For example, this parameter can be used to replace the standard [802.1Q] TPID value 0x88a8 with a value that is expected by the upstream device that supports any PB tagging, which uses any TPIDs like 0x9100 and 0x9200.

SubType	Length	Value
43.5.14.1	2	2-byte TPID value

B.3.15.2 Downstream TPID Translation

This TLV defines the new TPID value to be used for the outermost Tag of the frame before it is transmitted on the downstream interface. For example, this TLV can be used to replace any TPID value of Provider Bridge S-Tag (e.g., 0x9100) with a value that is expected by the downstream device that supports standard [802.1Q] tagging, which uses TPID 0x88a8.

SubType	Length	Value
43.5.14.2	2	2-byte TPID value

B.3.15.3 Upstream S-TPID Translation

This TLV defines the new TPID value to be used for the outermost S-Tag of the frame before it is transmitted on the NSI interface. For example, this TLV can be used to replace the standard [802.1Q] TPID value 0x88a8 with a value that is expected by the upstream device that supports any PB tagging, which uses any TPIDs like 0x9100 and 0x9200.

SubType	Length	Value
43.5.14.3	2	2-byte TPID value

B.3.15.4 Downstream S-TPID Translation

This TLV defines the new TPID value to be used for the outermost S-Tag of the frame before it is transmitted on the downstream interface. For example, this TLV can be used to replace any TPID value of PB S-Tag (e.g., 0x9100) with a value that is expected by the downstream device that supports standard [802.1Q] tagging, which uses TPID 0x88a8.

SubType	Length	Value
43.5.14.4	2	2-byte TPID value

B.3.15.5 Upstream B-TPID Translation

This TLV defines the new TPID value to be used for the outermost B-Tag of the frame before it is transmitted on the NSI interface. For example, this TLV can be used to replace the standard [802.1Q] TPID value 0x88a8 with a value that is expected by the upstream device that supports any PBB tagging.

SubType	Length	Value
43.5.14.5	2	2-byte TPID value

B.3.15.6 Downstream B-TPID Translation

This parameter defines the new TPID value to be used for the outermost B-Tag of the frame before it is transmitted on the downstream interface. For example, this parameter can be used to replace any TPID value of PBB B-Tag with a value that is expected by the downstream device that supports standard [802.1Q] tagging, which uses TPID 0x88a8.

SubType	Length	Value
43.5.14.6	2	2-byte TPID value

B.3.15.7 Upstream I-TPID Translation

This TLV defines the new TPID value to be used for the outermost I-Tag of the frame before it is transmitted on the NSI interface. For example, this TLV can be used to replace the standard [802.1Q] TPID value 0x88e7 with a value that is expected by the upstream device that supports PBB tagging.

SubType	Length	Value
43.5.14.7	2	2-byte TPID value

B.3.15.8 Downstream I-TPID Translation

This TLV defines the new TPID value to be used for the outermost I-Tag of the frame before it is transmitted on the downstream interface. For example, this TLV can be used to replace a TPID value of PBB I-Tag with a value that is expected by the downstream device that supports standard [802.1Q] tagging, which uses TPID 0x88e7.

SubType	Length	Value
43.5.14.8	2	2-byte TPID value

B.3.16 L2CP Processing

This parameter is used in DPoE specifications only. It is not used in DOCSIS L2VPN configuration.

The TLVs defined in this section are used for the configuration of L2CP in the DPoE Network. This TLV is applicable only to DPoE networks at this time. The L2CP Processing can only be applied in Encapsulation Mode.

SubType	Length	Value
43.5.15	n	

B.3.16.1 L2CP Tunnel Mode

This TLV defines the Tunnel Mode for L2CP.

SubType	Length	Value
43.5.15.1	1	L2CP Tunnel Mode 0 = Tunnel with L2PT 1 = Tunnel without L2PT (default) 2-255 = Reserved

B.3.16.2 L2CP D-MAC Address

This TLV defines the D-MAC address of the L2CP frames to be processed. Operators can use this TLV to provision the DA MAC address of the L2CP frames for L2CP processing

SubType	Length	Value
43.5.15.2	6	48-bit MAC Address of the L2CP DA MAC address

B.3.16.3 L2CP Overwriting D-MAC Address

This TLV defines the replacing D-MAC address for the L2CP frames in D-MAC Change mode. Operators can use this TLV to provision the overwriting D-MAC address for the L2CP frames.

SubType	Length	Value
43.5.15.3	6	48-bit MAC Address of the replacing D-MAC address

B.3.17 DAC Disable/Enable Configuration

This parameter is used in DPoE specifications only. It is not used in DOCSIS L2VPN configuration.

This TLV defines the Demarc Auto Configuration administrative configuration of the port (s) associated with a MEF service. Please refer to the DPoE Specifications for further detail [DPoE MEFv2.0] and [DPoE DEMARCV1.0].

This TLV can be configured with value 0 or 1 where 0 is for Disabling the DAC, and 1 is for Enabling the DAC. Without specifying this TLV, the DAC is disabled by default.

SubType	Length	Value
43.5.16	1	1 or 0

B.3.18 Pseudowire Class

This parameter is used in DPoE specifications only. It is not used in DOCSIS L2VPN configuration.

This STRING-based TLV defines the PW-Class to be used in creating the Pseudo Wire. The Pseudowire Class applies to the MPLS Peer, Pseudowire ID, Backup MPLS Peer, and Backup Pseudowire ID. If no MPLS Peer or Pseudowire ID is present in the CM configuration file, no Pseudo Wire class is needed.

SubType	Length	Value
43.5.18	1-32	Arbitrary

B.3.19 Service Delimiter

This parameter is used in DPoE specifications only. It is not used in DOCSIS L2VPN configuration.

This TLV defines a Service Delimiter, which consists of selector-bytes to be used by the Bridge Forwarder to associate the frames from a defined Tran-trail to the service instance or Pseudowire forwarder. The Service Delimiter TLV includes a list of Sub-TLVs that define the C-VID, S-VID, I-SID, and B-VID. See [DPoE ARCHv1.0] for more details.

SubType	Length	Value
43.5.19	n	

B.3.19.1 C-VID

This TLV defines [802.1Q] C-VID value to be used for service delimiting.

SubType	Length	Value
43.5.19.1	2	This TLV comprises an encoded bit map, featuring one field: C-VID, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	4 bits
C-VID	Encodes the C-VID field	12 bits

B.3.19.2 S-VID

This TLV defines [802.1Q] S-VID value to be used service delimiting.

SubType	Length	Value
43.5.19.2	2	This TLV comprises an encoded bit map, featuring one field: S-VID, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	4 bits
S-VID	Encodes the S-VID field	12 bits

B.3.19.3 I-SID

This TLV defines 24-bits [802.1Q] I-SID Backbone Service Instance Identifier value to be used for service delimiting.

SubType	Length	Value
43.5.19.3	3	24-bit value of [802.1Q] I-SID Backbone Service Instance Identifier

B.3.19.4 B-VID

This TLV defines [802.1Q] B-VID value to be used for service delimiting.

SubType	Length	Value
43.5.19.4	2	This TLV comprises an encoded bit map, featuring one field: B-VID, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	4 bits
B-VID	Encodes the B-VID field	12 bits

B.3.20 Virtual Switch Instance Encoding

This parameter is used in DPoE specifications only. It is not used in DOCSIS L2VPN configuration. This TLV defines a Virtual Switch Instance, which consists of series of TLVs that identify the configuration of the VSI on the CMTS. See [DPoE ARCHv1.0], [DPoE IPNEv1.0] and [DPoE MEFv2.0] for more details.

SubType	Length	Value
43.5.20	n	

B.3.20.1 VPLS Class

This STRING-based TLV is used to define standard operator VPLS behavior. See [DPoE ARCHv1.0], [DPoE IPNEv1.0] [DPoE MEFv2.0] for more details.

SubType	Length	Value
43.5.20.1	1-32	Arbitrary

B.3.20.2 E-Tree Role

This Optional TLV defines the role of the attachment circuit within the VSI. ROOT and LEAF are the only two valid values allowed for the E-Tree Role. See [DPoE ARCHv1.0], [DPoE IPNEv1.0] and [DPoE MEFv2.0] for more details.

SubType	Length	Value
43.5.20.2	1	0 – ROOT 1 – LEAF

B.3.20.3 E-Tree Root VID

This INTEGER-based TLV defines the appropriate egress or ingress based VID/frame association rules. See [DPoE ARCHv1.0], [DPoE IPNEv1.0] and [DPoE MEFv2.0] for more details.

SubType	Length	Value
43.5.20.3	2	This TLV comprises an encoded bit map, featuring one field: E-Tree ROOT VID, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	4 bits
E-Tree-Root-VID	Encodes the E-Tree-ROOT-VID field	12 bits

B.3.20.4 E-Tree Leaf VID

This INTEGER-based TLV defines the appropriate egress or ingress based VID/frame association rules. See [DPoE ARCHv1.0], [DPoE IPNEv1.0] and [DPoE MEFv2.0] for more details.

SubType	Length	Value
43.5.20.4	2	This TLV comprises an encoded bit map, featuring one field: E-Tree LEAF VID, as shown in the table below

Field name	Description	Size
Reserved	Reserved, ignored on reception	4 bits
E-Tree-Leaf-VID	Encodes the E-Tree-LEAF-VID field	12 bits

B.3.21 BGP Attribute sub TLV

The BGP Attribute is an Optional attribute. However, it is required when BGP auto-discovery is to be used by the CMTS for L2VPN discovery. There may be one or more BGP attributes present within a CM configuration file. See [DPoE ARCHv1.0], [DPoE IPNEv1.0] and [DPoE MEFv2.0] for more details.

SubType	Length	Value
43.5.21	N	

B.3.21.1 BGP VPNID

The BGP VPNID is a Mandatory sub-TLV of the BGP object. The BGP VPNID uniquely identifies an L2VPN to auto-discovery protocols that dynamically discover the remote L2VPN forwarders (i.e., PE routers) for MPLS Pseudowires. The BGP VPNID is unique across *all* CMTSs and PE routers in an administrative domain. The CMTS MUST use the value of the BGP VPNID to establish the L2VPN with remote PE routers. The BGP VPNID is assumed to be (or map to) the L2VPN Extended Community in BGP for the IETF L2VPN Auto-Discovery Protocol [RFC 6074]. A cable operator may use the BGP VPNID integer value in VPNID TLV as well.

SubType	Length	Value
43.5.21.1	4	An Integer that identifies a Layer 2 Virtual Private Network.

B.3.21.2 Route Distinguisher

The Route Distinguisher is an Optional sub-TLV of the BGP object used in DPoE. The purpose of the RD is defined within [RFC 4364]. The format of the Route Distinguisher is <Router ID>:<VPNID>, where the <Router ID> is the IPv4 address of the BGP router and <VPNID> is the value in TLV43.5.21.4. See [DPoE ARCHv1.0], [DPoE IPNEv1.0], and [DPoE MEFv2.0] for more details.

SubType	Length	Value
43.5.21.2	8	Integer

B.3.21.3 Route Target (Import)⁸⁸

The Route Target (import) is an Optional sub-TLV of the BGP object and should be used only if the default Route Target value (that is auto-calculated by the CMTS) is not good enough (e.g., E-TREE). The CMTS will use this sub-TLV to override the default RT value. The purpose of the RT is defined within [RFC 4364]. The Route Target (import) TLV is a set of import Route Targets, which, as a consequence, imply the possibility for multiple Route Target (import) TLVs to be present and associated with a single BGP object. The format of the RT is described in [RFC 4364] and it consists of a 2 byte Type field and a 6 byte Value field. Depending on the type the Value field can be encoded as: <2 byte ASN>:<4 byte value>, <IPv4 Address>:<2 byte value> or <4 byte ASN>:<2 byte value>.

SubType	Length	Value
43.5.21.3	Multiples of 8	Integer

B.3.21.4 Route Target (Export)⁸⁹

The Route Target (export) is an Optional sub-TLV of the BGP object and should be used only if the default Route Target value (that is auto-calculated by the CMTS) is not good enough (e.g., E-TREE). The CMTS will use this sub-TLV to override the default RT value. The purpose of the RT is defined within [RFC 4364]. The Route Target (export) TLV is a set of import Route Targets, which as a consequence implies the possibility for multiple Route

⁸⁸ Revised per L2VPN-N-13.1009-3 on 7/23/13 by PO.

⁸⁹ Revised per L2VPN-N-13.1009-3 on 7/23/13 by PO.

Target (export) TLVs to be present and associated with a single BGP object. The format of the Route Target is described in [RFC 4364] and it consists of a 2 byte Type field and a 6 byte Value field. Depending on the type the Value field can be encoded as: <2 byte ASN>:<4 byte value>, <IPv4 Address>:<2 byte value> or <4 byte ASN>:<2 byte value>.

SubType	Length	Value
43.5.21.4	Multiples of 8	Integer

B.3.21.5 CE-ID/VE-ID⁹⁰

This optional sub-TLV is used to configure CE-ID or VE-ID. For point-to-point, the value is either 4 (Ethernet Tagged Mode) or 5 (Ethernet Raw Mode). For multipoint, the value is 19 (VPLS).

SubType	Length	Value
43.5.21.5	2	CE-ID or VE-ID as determined by 43.5.2.4.3

B.3.22 VPN-SG Attribute sub TLV

This parameter is used in DPoE specifications only. It is not used in DOCSIS L2VPN configuration.

The SG attribute is critical for associating a service-flow with a serving-group as described in [DPoE MULPIv1.0] and [DPoE IPNEv1.0]. The value of this TLV points to a Serving Group identifier on the DPoE System. See [DPoE ARCHv1.0], [DPoE IPNEv1.0], and [DPoE MEFv2.0] for more details.

SubType	Length	Value
43.5.22	1...16	String

B.3.23 Pseudowire Signaling

This TLV defines the pseudowire signaling protocol to be used for MPLS Encapsulation.

SubType	Length	Value
43.5.23	1	0 : BGP 1: LDP 2: L2TPv3

B.3.24 L2VPN SOAM Subtype

This subtype is used to configure the SOAM functionality on the CM. A valid L2VPN Encoding may have multiple L2VPN SOAM Subtypes.

SubType	Length	Value
43.5.24	n	Various sub-TLVs

B.3.24.1 MEP Configuration

This subtype is used to configure the MEP functionality on the CM. The MEP is configured on a per-L2VPN basis.

SubType	Length	Value
43.5.24.1	n	Various sub-TLVs

⁹⁰ Revised per L2VPN-N-13.1009-3 on 7/23/13 by PO.

B.3.24.1.1 MD Level

This subtype is used to configure the Service OAM Maintenance Entity Group Level (MEG level) or Maintenance Domain Level (MD Level).

It accepts a value between 0 and 7, corresponding to the MD levels as defined in [MEF 30.1].

SubType	Length	Value
43.5.24.1.1	1	MD Level (0..7)

B.3.24.1.2 MD Name

This subtype is used to configure the Maintenance Domain Name of the Maintenance Association Identifier (MAID).

SubType	Length	Value
43.5.24.1.2	2..16	ASCII String

B.3.24.1.3 MA Name

This subtype is used to configure the Maintenance association Name of the Maintenance Association Identifier (MAID).

SubType	Length	Value
43.5.24.1.3	2..28	ASCII String

B.3.24.1.4 MEP ID⁹¹

This subtype is used to configure the Maintenance association Endpoint Identifier (MEP ID).

SubType	Length	Value
43.5.24.1.4	2	1..8191

B.3.24.2 Remote MEP Configuration

This subtype is used to configure remote MEPs with which the CM communicates. Multiple remote MEPs can be configured.

SubType	Length	Value
43.5.24.2	n	Various sub-TLVs

B.3.24.2.1 MD Level

This subtype is used to configure the Service OAM Maintenance Entity Group Level (MEG level) / Maintenance Domain Level (MD Level) for Remote Peers.

It accepts a value between 0 and 7, corresponding to the MD levels as defined in [MEF 30.1]. When the CM supports a single MEP, the CM ignores this TLV if present and instead uses its configured MD Level for peer MEPs.

SubType	Length	Value
43.5.24.2.1	1	MD Level (0..7)

⁹¹ Revised per L2VPN-N-13.1009-3 on 7/23/13 by PO.

B.3.24.2.2 MD Name

This subtype is used to configure the Maintenance Domain Name of the Maintenance Association Identifier (MAID).

SubType	Length	Value
43.5.24.2.2	2..16	ASCII String

B.3.24.2.3 MA Name

This subtype is used to configure the Maintenance association Name of the Maintenance Association Identifier (MAID).

SubType	Length	Value
43.5.24.2.3	2..28	ASCII String

B.3.24.2.4 MEP ID⁹²

This subtype is used to configure the Maintenance association Endpoint Identifier (MEP ID).

SubType	Length	Value
43.5.24.2.4	2	1..8191

B.3.24.3 Fault Management Configuration

This subtype is used to configure the Fault Management functionality.

SubType	Length	Value
43.5.24.3	n	Various sub-TLVs

B.3.24.3.1 Continuity Check Messages

This subtype is used to enable and configure the transmission periodicity of CCM.

SubType	Length	Value
43.5.24.3.1	1	0: CCM is disabled. 1: Transmission periodicity of 10 seconds 2: Transmission periodicity of 1 second

B.3.24.3.2 Loopback Function⁹³

This subtype is used to enable the Loopback functionality.

SubType	Length	Value
43.5.24.3.2	1	0 : Disabled 1: Loopback Responder function enabled (e.g., receive LBM and reply with LBR)

⁹² Revised per L2VPN-N-13.1120-2 on 11/7/13 by PO.

⁹³ Revised per L2VPN-N-14.1135-5 on 3/31/14 by JB.

B.3.24.3.3 Linktrace Function⁹⁴

This subtype is used to enable the Linktrace functionality.

SubType	Length	Value
43.5.24.3.3	1	0 : Disabled 1: Linktrace Responder function enabled (e.g., receive LTM and reply with LTR)

B.3.24.4 Performance Management Configuration

This subtype is used to configure the Performance Management functionality.

SubType	Length	Value
43.5.24.4	n	Various sub-TLVs

B.3.24.4.1 Frame Delay Measurement

This subtype is used to configure the Frame Delay measurement functionality.

SubType	Length	Value
43.5.24.4.1	n	Various sub-TLVs

B.3.24.4.1.1 Frame Delay Measurement Enable⁹⁵

This subtype is used to enable Frame Delay measurement.

SubType	Length	Value
43.5.24.4.1.1	1	0: Disabled 1: Frame Delay Responder function enabled (e.g., receive DMM and reply with DMM)

B.3.24.4.1.2 Frame Delay Measurement One-way-Two-way

This subtype is used to configure the Frame Delay measurement as one-way or two-way.

SubType	Length	Value
43.5.24.4.1.2	1	0: One-way 1: Two-way

B.3.24.4.1.3 Frame Delay Measurement Transmission Periodicity

This subtype is used to configure the transmission periodicity of the Frame Delay measurement.

SubType	Length	Value
43.5.24.4.1.3	2	Integer times in milliseconds. If value is = 0, then FD test is to be run once.

B.3.24.4.2 Frame Loss Measurement

This subtype is used to configure Frame Loss measurement.

⁹⁴ Revised per L2VPN-N-14.1135-5 on 3/31/14 by JB.

⁹⁵ Revised per L2VPN-N-14.1135-5 on 3/31/14 by JB.

SubType	Length	Value
43.5.24.4.2	n	Various sub-TLVs

B.3.24.4.2.1 Frame Loss Measurement Enable⁹⁶

This subtype is used to enable Frame Loss measurement.

SubType	Length	Value
43.5.24.4.2.1	1	0: Disabled 1: Frame Loss Responder function enabled (e.g., receive SLM and reply with SLR)

B.3.24.4.2.2 Frame Loss Measurement Transmission Periodicity

This subtype is used to configure the transmission periodicity of the Frame Loss measurement.

SubType	Length	Value
43.5.24.4.2.2	1	Integer times in milliseconds. If value is = 0, then FLM test is to be run once.

B.3.25 Network Timing Profile Reference

This parameter is used in DPoE specifications only. It is not used in DOCSIS L2VPN configuration.

The Network Timing Profile Reference is used to associate an L2VPN Service Flow to a Network Timing Profile Name in the CM configuration file. A valid Network Timing Profile subtype encoding contains one instance of this subtype.

Type	Length	Value
43.5.25	2	Network Timing Profile Reference

B.3.26 L2VPN DSID

This parameter is only used for DOCSIS 3.0 L2VPN systems.

Type	Length	Value
43.5.26	3	DSID (1-1048575)

The CMTS includes this TLV when Multicast DSID Forwarding is enabled, and the CMTS does not otherwise sequence L2VPN Multicast traffic to the CM for downstream bonding. The CMTS may label L2VPN multicast traffic with an L2VPN DSID. The CM associates no forwarding attributes with the L2VPN DSID.

⁹⁶ Revised per L2VPN-N-14.1135-5 on 3/31/14 by JB.

B.4 Confirmation Codes⁹⁷

This section defines new confirmation codes for L2VPN operation. It extends the list of confirmation codes in [MULPIv3.0].⁹⁸

Additional Confirmation Codes defined for the DOCSIS L2VPN feature include:

- reject-VLAN-ID-in-use (100): indicates that an IEEE802.1Q VLAN-ID requested for the NSI encapsulation of L2VPN traffic is already assigned for use by non-L2VPN traffic. See Section 6.2.5.
- reject-multipoint-L2VPN(101): indicates that Multipoint L2VPN forwarding mode is not supported and a CM is attempting to configure more than one L2VPN attachment circuit to the same L2VPN. See Section 6.2.5.
- reject-multipoint-NSI(102): indicates that a multipoint forwarding L2VPN contained multiple L2VPN encodings with different NSI encapsulation values.

B.5 L2VPN Error Encoding

This encoding provides additional information from the CM when it rejects an L2VPN Encoding signaled by the CMTS. The CM MUST include an L2VPN Error Encoding in its MAC management response when it rejects an L2VPN Encoding in a REG-RSP, DSA-REQ, DSA-RSP, DSC-REQ or DSC-RSP.

GEI Type	Length	Value
43.5.254	N	L2VPN Error Encoding, consisting of exactly one L2VPN Errored Parameter encoding, exactly one L2VPN Error Code encoding, and zero or one L2VPN Error Message encoding.

B.5.1 L2VPN Errored Parameter

This parameter provides a sequence of Type and Subtypes that identify the location and subtype of the L2VPN Encoding that is rejected. A valid L2VPN Error Encoding contains exactly one L2VPN Errored Parameter type string.

GEI Type	Length	Value
43.5.254.1	N	Sequence of Type and Subtypes

The type/subtype sequence starts at the top level of TLV encodings of the MAC Management Message that included the L2VPN Encoding. This sequence depends on the location of the L2VPN encoding, as described in Section 6.2. In particular:

- An L2VPN Error Parameter string for a top-level L2VPN Encoding starts with two bytes for the GEI Type code for the L2VPN Encoding, or (43.5);
- An L2VPN Error Parameter string for an Upstream Service Flow Encoding starts with the type code for that encoding (24) followed by the L2VPN Encoding GEI Type, or (24.43.5);
- An L2VPN Error Parameter string for a Downstream Packet Classification Configuration Setting starts with the type for that encoding (23) followed by the L2VPN Encoding GEI Type, or (23.43.5);
- An L2VPN Error Parameter string for an Upstream Packet Classification Configuration Setting starts with the type for that encoding (22) followed by the L2VPN Encoding GEI Type, or (22.43.5).

⁹⁷ Section modified per L2VPN-N-14.1186-1 on 1/15/15 by PO.

⁹⁸ Sentence and bullets in next paragraph updated per L2VPN-N-07.0460-1 on 7/25/07 by KN.

If the entire L2VPN Encoding is rejected, the CM MAY include in the L2VPN Error Parameter type string only the two or three bytes that identify the location of a full L2VPN Encoding. If the reason for rejection is due to a particular subtype of the L2VPN Encoding, the CM SHOULD include additional bytes in the L2VPN Error Parameter type string to identify the particular subtype of the L2VPN Encoding that it rejected. One reason for rejecting an entire L2VPN encoding is that the maximum number of L2VPNs supported by the CM has been exceeded. One reason for rejecting a particular subtype, e.g., the L2VPN SA-Descriptor Subtype Encoding, is that the number of SAIDs supported by the CM has been exceeded.

B.5.2 L2VPN Error Code

This parameter provides a confirmation code as defined from [RFI 2.0] to identify a reason why an L2VPN Encoding or subtype was rejected. A valid L2VPN Error Encoding contains exactly one L2VPN Confirmation code.

GEI Type	Length	Value
43.5.254.2	1	Confirmation Code

B.5.3 L2VPN Error Message

This parameter, if present, provides a message for display on the CMTS console log for the reason for the rejection. A CM SHOULD include this parameter in an L2VPN Error Encoding. A valid L2VPN Error Encoding contains zero or one L2VPN Error Message subtypes.

GEI Type	Length	Value
43.5.254.3	N	Zero-terminated string of ASCII characters.

B.6 CM Interface Mask Classification Criteria

This specification defines a generic mechanism for classifying upstream and downstream traffic based on the ingress or intended egress logical interface ports on the CM.

In an Upstream Packet Classifier Encoding (type 22), the CM Interface Mask Subtype defines a rule criteria for matching the ingress interface of an L2PDU.

In a Downstream Packet Classifier Encoding (type 23), the CM Interface Mask Subtype defines a rule criteria for matching a unicast downstream destination MAC address. In either case, the CMIM Subtype encoding within a Packet Classifier Encoding applies to both L2VPN and non-L2VPN traffic.

Each bit of CMIM corresponds to a logical bridge port interface of a MAC layer 2 bridge implemented in the eCM of a cable modem. The parameter is encoded as the octet string of the Basic Encoding Rules encoding of an SNMP BITS bit string. Bit position K in the BITS encoding corresponds to eDOCSIS MAC bridge interface K. By convention, bit position 0 corresponds to the eCM's self host interface (i.e., the CM's IP stack). The eCM self MAC address is signaled as if it were on a bridge port interface ifIndex of zero (0), even though no such interface actually exists.

SubType	Length	Value
[22/23].13	N	<p>SNMP BITS -encoded bit map with bit position K representing eCM logical interface index value K. Bit position 0 represents the eCM self host itself. Bit position 0 is the most significant bit of the first octet. The Embedded DOCSIS specification [eDOCSIS] defines the interface index assignments. For information purposes, current assignments include:</p> <p>Bit 0 (0x80): eCM self host interface</p> <p>Bit 1 (0x40): primary CPE Interface (also eRouter) ⁹⁹</p> <p>Bit 2 (0x20) RF interface</p> <p>Bits 3,4 reserved</p> <p>Bits 5..15 (0x07 FF) Other CPE Interfaces</p> <p>Bits 16-31, Logical CPE Interfaces for eSAFE hosts. Current assignments include:</p> <p>Bit 16 (0x00 00 80) PacketCable-EMTA</p> <p>Bit 17 (0x00 00 40) eSTB-IP</p> <p>Bit 18 (0x00 00 20) eSTB-DSG</p> <p>Bits 19..31 (0x00 00 1F FF) Other eSAFE interfaces</p>

In an Upstream Classifier Encoding, a CM MUST silently ignore bit positions for unimplemented interfaces. For example, an upstream CMIM classifier criteria intended to match only external CPE interfaces of a CM has a CMIM mask value setting bits 1 and 5-15, i.e., an encoding of 0x47 FF.

In a Downstream Classifier Encoding, that includes a CMIM criteria, the CMTS checks the destination MAC address to determine whether it is the CM's self host MAC address or a recognized eSAFE host MAC address. Any other unicast MAC address is considered to be a CPE MAC address. The CMTS does not know on what particular CPE interface the CM has learned a CPE MAC address. The CMTS considers only bit 1 of the CMIM to match a CPE MAC address in a Downstream Packet Classifier Encoding. The maximum number of eSAFE destination MAC addresses recognized by a CMTS is vendor specific.

B.7 L2VPN MAC Aging Encoding ¹⁰⁰

The L2VPN MAC aging parameter is intended to allow CMs to provide services compatible with MEF specifications. In some cases such as when offering MEF-certified service, it is necessary for the CM to replace MAC addresses in the CM bridging table with newly discovered addresses when the table is full. This parameter enables such CMs to age out the MAC address of the device that has not transmitted a packet destined for the upstream RF interface for the longest amount of time when the CM bridging table is full.

Top-Level TLV Encoding for L2VPN MAC Aging Control:

Type	Length	Value
65	N	Composite value

The encapsulated fields are described below.

⁹⁹ Revised per L2VPN-N-12.1065-11 on 9/19/12 by CP.

¹⁰⁰ Section added per L2VPN-N-08.0625-2 on 4/16/08 by KN.

B.7.1 L2VPN MAC Aging Mode

Type	Length	Value
65.1	1	L2VPN MAC Aging Mode 0 = Disable (default) 1 = Enable 2-255 = Reserved.

This setting is found in the configuration file.

The L2VPN MAC Aging Control encoding is not intended to be forwarded by the CM to the CMTS in the Registration Request message. As such, it is not expected to be included in the E-MIC Bitmap for DOCSIS 3.0 devices or in the CMTS MIC calculation for DOCSIS 2.0 devices.

If the L2VPN MAC Aging Control encoding is omitted or the L2VPN MAC Aging Mode bit is zero, then the CM uses the MAC address acquisition requirements in [DOCSIS MAC]. This means that the CM will not replace previously acquired addresses with newly discovered CPE MAC addresses.

- If the L2VPN MAC Aging Mode is set to "1", the CM MUST use the L2VPN MAC Aging Mode, which supersedes only the CM MAC address acquisition and filtering rules defined in [DOCSIS MAC].¹⁰¹ If the CM has acquired the maximum number of CPE MAC addresses allowed by its Max CPE value and it discovers a new source MAC address sending upstream traffic, then:
 - The CM MUST preserve statically provisioned MAC addresses in its bridging table.
 - The CM MUST preserve eSAFE MAC addresses in its bridging table.
 - The CM MUST update its bridging table to replace the MAC address of a device that has not transmitted a packet destined for the modem's upstream RF interface for the longest amount of time with the newly discovered CPE MAC address.
 - The MAC Address Learning Control Encoding (TLV 69) as defined in [MULPIv3.0], if configured, applies to a CM when the L2VPN MAC Aging Mode is set to either "0" or "1".¹⁰²

¹⁰¹ Sentence revised per L2VPN-N-12.1065-11 on 9/19/12 by CP.

¹⁰² Bullet added per L2VPN-N-12.1065-11 on 9/19/12 by CP.

Appendix I Example L2VPN Encodings (Informative)

The L2VPN Encoding is always encapsulated using a General Extension Information (GEI) encoding, which uses the type code 43 with the reserved Vendor ID of 0xFFFFF.

I.1 Point-to-Point Example

This section describes L2VPN Encodings for three CMs performing point-to-point L2VPN forwarding of all traffic on their default upstream service flow. Two of the CMs are externally bridged to the same enterprise (L2VPN ID 0234560001); one of the CMs is bridged to a separate enterprise (L2VPN ID 0234560002). The example is depicted below:

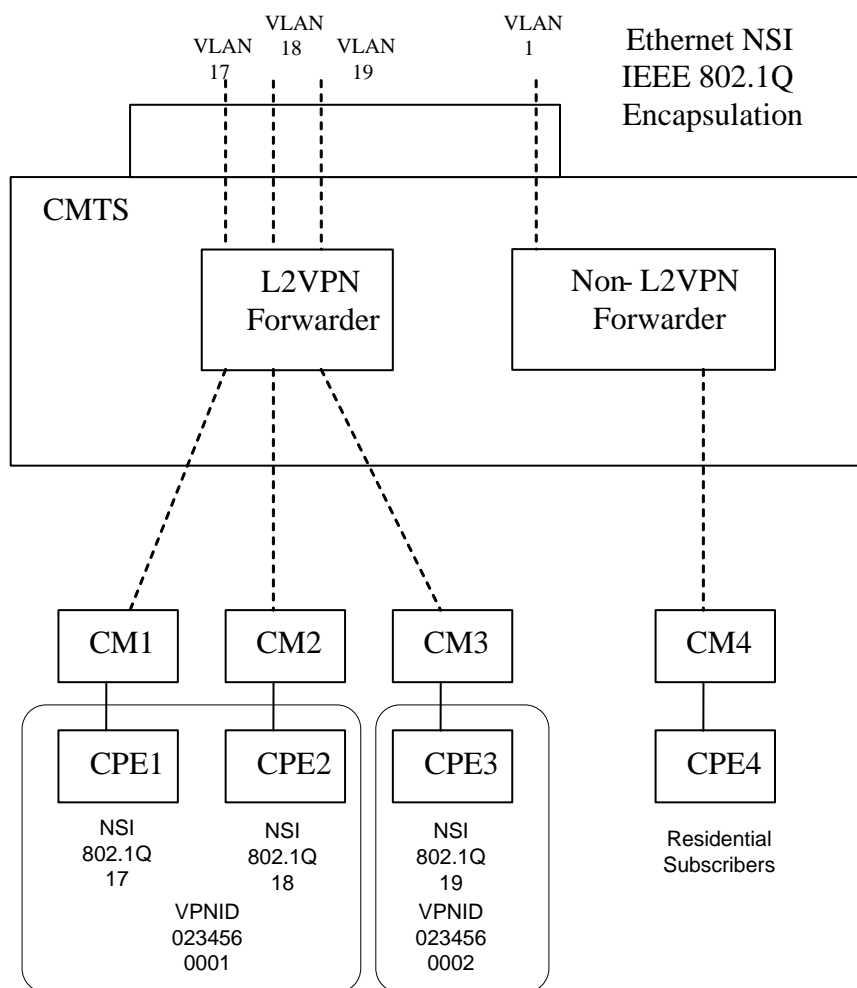


Figure I-1 - Point-To-Point L2VPN Traffic Forwarding Example

Table I-1 - Point-to-Point CM1 L2VPN Encoding¹⁰³

Point-to-Point CM1 Configuration File				
43				Per-CM L2VPN Encoding
20				Overall length
	08 03 FFFFFFFF			Vendor ID : 0xFFFFFFFF for GEI
	05			GEI 43.5 for L2VPN Encoding
	13			Length of GEI.5 Subtype
		01 05 x0234560001		VPNID Subtype
		02		NSI Encapsulation Subtype
		04		Length of GEI.5.2 Subtype
			02	IEEE 802.1Q Format Subtype
			02	Length of GEI.5.2.2 Subtype
			0x0011	VLAN ID 17
24				Upstream Service Flow Encoding
19				Length
	6			QOS Param Set Type Subtype
	1			
		0x07		
	43			Vendor-Specific Subtype:
	14			Overall length
		08 03 FFFFFFFF		Vendor ID for GEI
		05		GEI 43.5 for L2VPN Encoding
		7		Length of GEI.5 Subtype
			01 05 x0234560001	VPNID Subtype
45				DUT Filtering:
03				Overall Length
	01 01 01			DUT Filtering enabled

Table I-2 - Point-to-Point CM2 L2VPN Encoding¹⁰⁴

Point-to-Point CM2 Configuration File				
43				Per-CM L2VPN Encoding
20				Overall length
	08 03 FFFFFFFF			Vendor ID : 0xFFFFFFFF for GEI
	05			GEI 43.5 for L2VPN Encoding
	13			Length of GEI.5 Subtype
		01 05 x0234560001		VPNID Subtype
		02		NSI Encapsulation Subtype
		04		Length of GEI.5.2 Subtype
			02	IEEE 802.1Q Format Subtype
			02	Length of GEI.5.2.2 Subtype

¹⁰³ Table modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.¹⁰⁴ Table modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.

Point-to-Point CM2 Configuration File				
			0x0012	VLAN ID 18
24				Upstream Service Flow Encoding
19				Length
	6			QOS Param Set Type Subtype
	1			
		0x07		
	43			Vendor-Specific Subtype:
	14			Overall length
		08 03 FFFFFFFF		Vendor ID for GEI
		05		GEI 43.5 for L2VPN Encoding
		7		Length of GEI.5 Subtype
			01 05 x0234560001	VPNID Subtype
45				DUT Filtering:
03				Overall Length
	01 01 01			DUT Filtering enabled

CPE2 is externally bridged to the same L2VPN as CPE1 (VPNID x0234560001), but all L2VPN forwarding for CPE2 occurs on the NSI IEE 802.1Q VLAN ID 18.

Table I-3 - Point-to-Point CM3 L2VPN Encoding¹⁰⁵

Point-to-Point CM3 Configuration File				
43				Per-CM L2VPN Encoding
20				Overall length
	08 03 FFFFFFFF			Vendor ID : 0xFFFFFFFF for GEI
	05			GEI 43.5 for L2VPN Encoding
	13			Length of GEI.5 Subtype
		01 05 x0234560002		VPNID Subtype
		02		NSI Encapsulation Subtype
		04		Length of GEI.5.2 Subtype
			02	IEEE 802.1Q Format Subtype
			02	Length of GEI.5.2.2 Subtype
			0x0013	VLAN ID 19
24				Upstream Service Flow Encoding
19				Length
	6			QOS Param Set Type Subtype
	1			
		0x07		
	43			Vendor-Specific Subtype:
	14			Overall length
		08 03 FFFFFFFF		Vendor ID for GEI

¹⁰⁵ Table modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.

Point-to-Point CM3 Configuration File				
		05		GEI 43.5 for L2VPN Encoding
		7		Length of GEI.5 Subtype
			01 05 x0234560002	VPNID Subtype
45				DUT Filtering:
03				Overall Length
	01 01 01			DUT Filtering enabled

I.2 Multipoint Example

This section provides an example of L2VPN encodings for Multipoint forwarding, as depicted below. For Multipoint forwarding, the NSI encapsulation for an L2VPN may be configured in either of two ways:

- As CMTS vendor specific configuration; or
- In the CM configuration file of one or more of the CMs in the L2VPN.

In the example below, the NSI encapsulation for each L2VPN appears in the CM configuration file for all CMs.

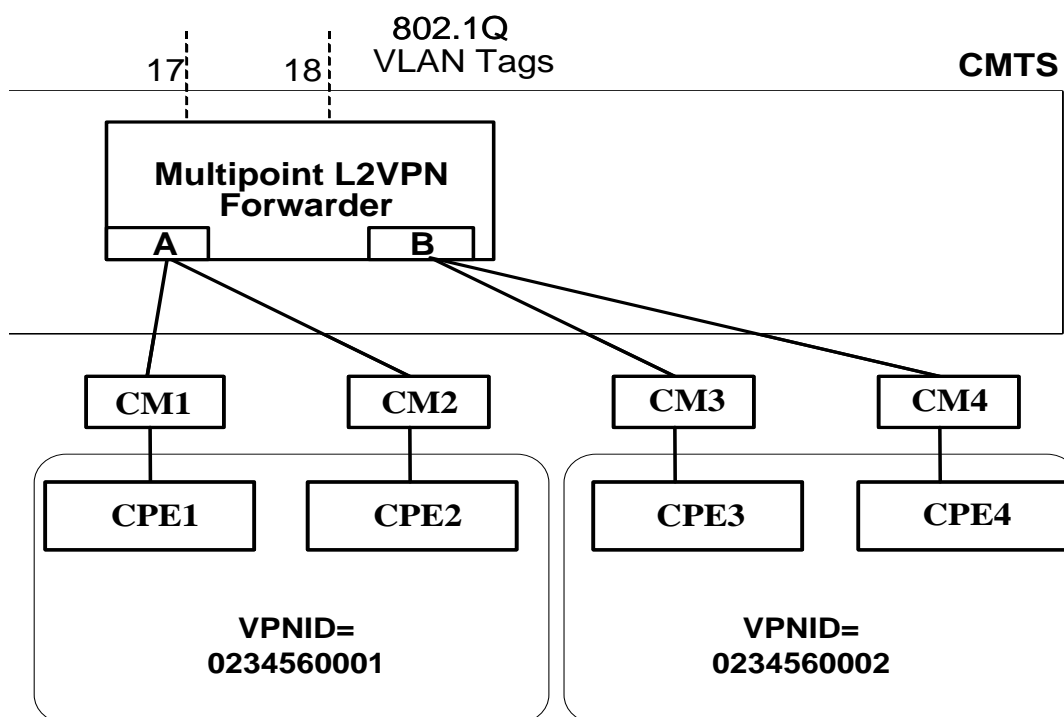


Figure I-2 - Multipoint L2VPN Forwarding of Traffic Example

Table I-4 - Multipoint CM1 L2VPN Encoding¹⁰⁶

Multipoint CM1 Config File				
43				Per-CM L2VPN Encoding
20				Overall length
	08 03 FFFFFFFF			Vendor ID : 0xFFFFFFFF for GEI
	05			GEI 43.5 for L2VPN Encoding
	13			Length of GEI.5 Subtype
		01 05 x0234560001		VPNID Subtype
		02		NSI Encapsulation Subtype
		04		Length of GEI.5.2 Subtype
			02	IEEE 802.1Q Format Subtype
			02	Length of GEI.5.2.2 Subtype
			0x0011	VLAN ID 17
24				Upstream Service Flow Encoding
19				Length
	6			QOS Param Set Type Subtype
	1			
		0x07		
	43			Vendor-Specific Subtype:
	14			Overall length
		08 03 FFFFFFFF		Vendor ID for GEI
		05		GEI 43.5 for L2VPN Encoding
		7		Length of GEI.5 Subtype
			01 05 x0234560001	VPNID Subtype
45				DUT Filtering:
03				Overall Length
	01 01 01			DUT Filtering enabled

Table I-5 - Multipoint CM2 L2VPN Encoding¹⁰⁷

Multipoint CM2 Config File				
43				Per-CM L2VPN Encoding
20				Overall length
	08 03 FFFFFFFF			Vendor ID : 0xFFFFFFFF for GEI
	05			GEI 43.5 for L2VPN Encoding
	13			Length of GEI.5 Subtype
		01 05 x0234560001		VPNID Subtype
		02		NSI Encapsulation Subtype
		04		Length of GEI.5.2 Subtype
			02	IEEE 802.1Q Format Subtype
			02	Length of GEI.5.2.2 Subtype
			0x0011	VLAN ID 17

¹⁰⁶ Table modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.¹⁰⁷ Table modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.

Multipoint CM2 Config File				
24				Upstream Service Flow Encoding
19				Length
	6			QOS Param Set Type Subtype
	1			
		0x07		
	43			Vendor-Specific Subtype:
	14			Overall length
		08 03 FFFFFFFF		Vendor ID for GEI
		05		GEI 43.5 for L2VPN Encoding
		7		Length of GEI.5 Subtype
			01 05 x0234560001	VPNID Subtype
45				DUT Filtering:
03				Overall Length
	01 01 01			DUT Filtering enabled

NOTE: The L2VPN Encodings for Multipoint CM2 are exactly the same as for CM1.

Table I-6 - Multipoint CM3 L2VPN Encoding¹⁰⁸

Multipoint CM3 Config File				
43				Per-CM L2VPN Encoding
20				Overall length
	08 03 FFFFFFFF			Vendor ID : 0xFFFFFFFF for GEI
	05			GEI 43.5 for L2VPN Encoding
	13			Length of GEI.5 Subtype
		01 05 x0234560002		VPNID Subtype
		02		NSI Encapsulation Subtype
		04		Length of GEI.5.2 Subtype
			02	IEEE 802.1Q Format Subtype
			02	Length of GEI.5.2.2 Subtype
			0x0012	VLAN ID 18
24				Upstream Service Flow Encoding
19				Length
	6			QOS Param Set Type Subtype
	1			
		0x07		
	43			Vendor-Specific Subtype:
	14			Overall length
		08 03 FFFFFFFF		Vendor ID for GEI
		05		GEI 43.5 for L2VPN Encoding
		7		Length of GEI.5 Subtype
			01 05 x0234560002	VPNID Subtype
45				DUT Filtering:

¹⁰⁸ Table modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.

Multipoint CM3 Config File				
03				Overall Length
	01 01 01			DUT Filtering enabled

Table I-7 - Multipoint CM4 L2VPN Encoding¹⁰⁹

Multipoint CM4 Config File				
43				Per-CM L2VPN Encoding
20				Overall length
	08 03 FFFFFFFF			Vendor ID : 0xFFFFFFFF for GEI
	05			GEI 43.5 for L2VPN Encoding
	13			Length of GEI.5 Subtype
		01 05 x0234560002		VPNID Subtype
		02		NSI Encapsulation Subtype
		04		Length of GEI.5.2 Subtype
			02	IEEE 802.1Q Format Subtype
			02	Length of GEI.5.2.2 Subtype
			0x0012	VLAN ID 18
24				Upstream Service Flow Encoding
19				Length
	6			QOS Param Set Type Subtype
	1			
		0x07		
	43			Vendor-Specific Subtype:
	14			Overall length
		08 03 FFFFFFFF		Vendor ID for GEI
		05		GEI 43.5 for L2VPN Encoding
		7		Length of GEI.5 Subtype
			01 05 x0234560002	VPNID Subtype
45				DUT Filtering:
03				Overall Length
	01 01 01			DUT Filtering enabled

NOTE: The L2VPN Encoding for Multipoint CM4 is the same as for CM3.

I.3 Upstream L2VPN Classifier Example

This example shows classifying upstream traffic from a specific CPE1 onto an upstream L2VPN service flow, where all other CPEs attached to the CM forward to the non-L2VPN forwarder, as depicted below.

Table I-8 - Upstream L2VPN Classifier Encoding¹¹⁰

Upstream L2VPN Classifier Cable Modem Config File				
43				Per-CM L2VPN Encoding
20				Overall length

¹⁰⁹ Table modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.

¹¹⁰ Table modified per L2VPN-N-07.0538-2 on 11/9/07 by KN.

Upstream L2VPN Classifier Cable Modem Config File				
	08 03 FFFFFFFF			Vendor ID: 0xFFFFFFFF for GEI
	05			GEI 43.5 for L2VPN Encoding
	13			Length of GEI.5 Subtype
		01 05 x0234560003		VPNID Subtype
		02		NSI Encapsulation Subtype
		04		Length of GEI.5.2 Subtype
			02	IEEE 802.1Q Format Subtype
			02	Length of GEI.5.2.2 Subtype
			0x0019	VLAN ID 25
24				Default Upstream Service Flow Encoding
03				Length
	06 01 07			QOS Param Set Type Subtype
24				L2VPN Upstream Service Flow Encoding
23				Length
	06 01 07			QOS Param Set Type Subtype
	01 02 0001			Service Flow Reference 0001
	43			Vendor-Specific Subtype:
	14			Overall length
		08 03 FFFFFFFF		Vendor ID for GEI
		05		GEI 43.5 for L2VPN Encoding
		7		Length of GEI.5 Subtype
			01 05 x0234560003	VPNID Subtype
22				Upstream Classifier Encoding
14				Length
	03 02 0001			Service Flow Reference to 0001
	10			Ethernet/LLC Packet Classification
		02		Source MAC Address
		6		Length
			x0001020000AA	MAC Address of CPE1
45				DUT Filtering:
03				Overall Length
	01 01 01			DUT Filtering enabled

Appendix II IEEE 802.1Q Encapsulation (Informative)

This appendix provides background information on the format of IEEE 802.1Q tags on Ethernet side NSI interfaces. This is the standard mechanism for indicating the VLAN of a bridged packet on an Ethernet interface. A CMTS compliant with this specification is required to support recognition of IEEE 802.1Q encapsulation on an Ethernet interface when configured to do so.

Because the CMTS interprets the VLAN ID of the outermost 802.1Q tag of a packet coming into an NSI, the tag is called a service-delimiting tag.

The L2VPN Forwarder *strips* the service-delimiting IEEE 802.1Q tag from an Ethernet packet when forwarding it downstream, and *inserts* the service-delimiting IEEE 802.1Q tag when forwarding packets upstream. The particular VLAN to which an L2VPN packet belongs is explicitly indicated on an 802.1Q-encapsulated Ethernet interface, and is always implied when forwarded on the DOCSIS RF MAC interface.

This stripping and inserting of IEEE 802.1Q tags is depicted below:

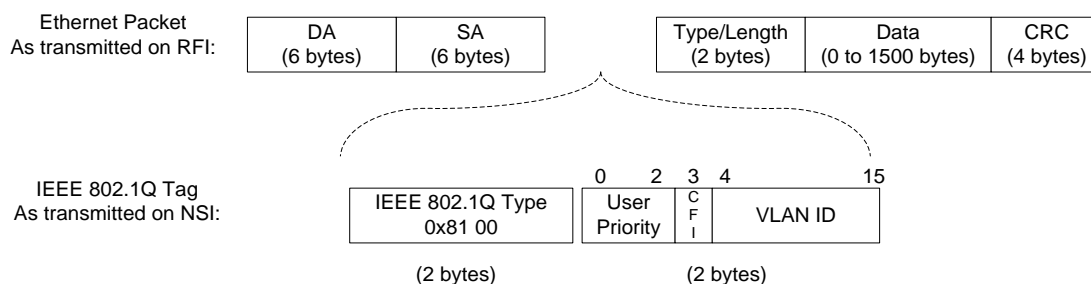


Figure II-1 - Ethernet 802.1Q Tags

An Ethernet packet is tagged with an IEEE 802.1Q tag by inserting four bytes between its original Source Address (SA) and original Length/Type field. The two-byte Ethernet Type code 0x8100 indicates that a 16-bit IEEE 802.1Q Tag follows. The tag value consists of a 3 bit user priority field in the most significant 3 bits, a Canonical Format Indicator (CFI) bit, and a 12-bit VLAN ID in the least significant bits. Operation of the CFI bit is defined by the IEEE, and is zero for Ethernet MAC addresses. All multi-byte fields are transmitted most significant byte first.

The User Priority field indicates a traffic forwarding priority in the range 0..7, with higher values indicating higher priority.

This specification permits, but does not require, the CMTS to use NSI port encapsulations other than IEEE 802.1Q to signal the L2VPN or attachment circuit for an L2VPN-forwarded packet. The particular NSI encapsulation used for L2VPN forwarding is intended to be configured in the NSI Encapsulation Subtype of an L2VPN Encoding.

Appendix III Embedded VLAN CM Bridging Model (Informative)

This appendix proposes an Embedded VLAN model for CM internal packet bridge forwarding for consideration by the DOCSIS community. It is not currently a requirement for L2VPN certification on a CM.

The L2VPN specification uses the concept of a CM Interface Mask (CMIM) to define the set of internal and external bridge interfaces to which the CM may bridge downstream traffic. The CMIM, for example, defines the broadcast domain of both individual and group MAC-addressed downstream traffic. A broadcast domain is one interpretation of a virtual LAN, so in effect, the CMIM is defining an internal VLAN of the internal and external ports to which DUT and L2VPN traffic are forwarded.

The Embedded VLAN model expands the MAC Bridge of the eDOCSIS model to become a VLAN-capable MAC bridge with separate embedded VLAN (eVLAN) MAC forwarding domains. By using the concept of an eVLAN, the CM is able to isolate the eCM and eSAFE hosts from the MAC broadcast domains of customer private L2VPNs.

The following Figure III–1 depicts the Embedded VLAN Model for L2VPN-compliant CMs.

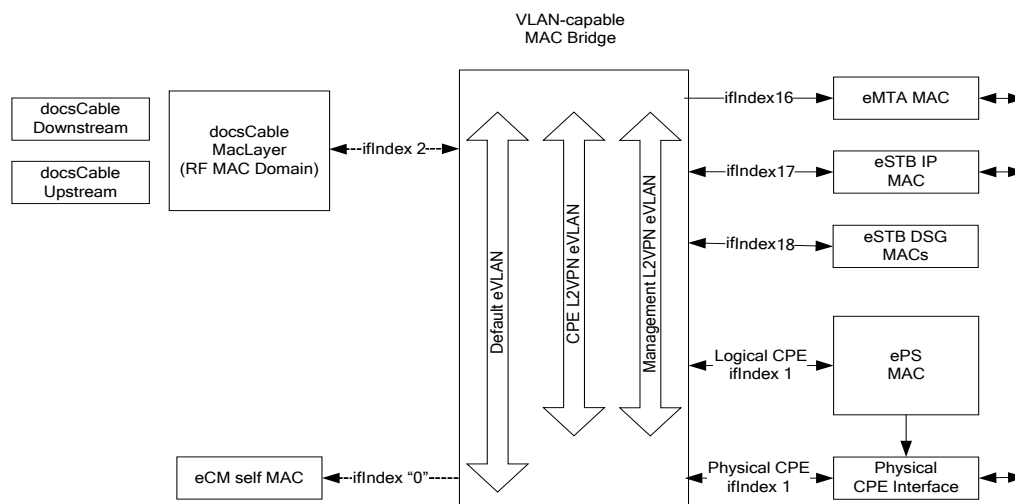


Figure III–1 - L2VPN Embedded VLAN (eVLAN) Model

The MAC Bridge of an embedded CM is considered to have a bridge port interface to the RF MAC Domain interface as ifIndex 2, and a Primary CPE bridge port interface at ifIndex 1. DOCSIS defines the operation of CPE forwarding by a residential CM as a layer-2 MAC bridging function between the RF interface and the CPE interface. In eDOCSIS, the eCM's own MAC address (its self MAC) is considered internal to the MAC bridge, and reachable by all bridge port interfaces.

The eDOCSIS specification defines an embedded Service/Application Functional Entity (eSAFE) as an entity co-located with an embedded Cable Modem (eCM) that contains its own MAC and IP address [eDOCSIS]. Currently defined eSAFEs include:

- PacketCable Embedded MTA (eMTA) host
- eRouter ¹¹¹
- ESTB Embedded Set Top Box

¹¹¹ Revised per L2VPN-N-12.1065-11 on 9/19/12 by CP.

Each of these eSAFE devices is considered to have a separate Logical CPE Interface to the MAC bridge, and is assigned a separate Interface Index (ifIndex) for management and control purposes. In the eDOCSIS architecture, the eCM's MAC Bridge is assumed to implement a single forwarding database, associating MAC addresses to each Logical CPE interface, and forwarding Layer 2 Protocol Data Units (L2PDUs) between all ports of the MAC bridge, according to the Destination MAC (DMAC) address of the L2PDU. The RF Interface, CPE interface(s), eCM MAC, and all eSAFE MACs are considered to be in the same layer 2 MAC broadcast domain, (i.e., on a single LAN).

The L2VPN specification will expand this architecture by introducing the concept of embedded VLANs (eVLANs) within the eCM's MAC bridge, where eVLANs have different sets of logical CPE ports. In order to control access to the eCM's self MAC address, (e.g., to isolate it from customer L2VPN access), the eCM MAC is considered to reside on a self-bridge port interface.

The L2VPN architecture introduces the concept of a Cable Modem Interface Mask (CMIM), with a bit position for each logical bridge port in the eCM's eVLAN-capable MAC Bridge. Each eVLAN in the MAC Bridge contains a CMIM value that represents which logical bridge ports belong to the eVLAN. The CMIM is represented as an SNMP BITS object encoding, where bit position K corresponds to bridge port interface ifIndex K. In a CMIM mask, the logical Self bridge port is assigned bit position 0 (i.e., as if it had ifIndex value 0). No ifStack entry is created for the Self bridge interface, because zero is an invalid value for an ifIndex value.

In the eCM's MAC Bridge, all non-L2VPN forwarding is considered to be bridged on a Default eVLAN that has a CMIM, with all interface bits set to 1. This corresponds to the normal, single-LAN MAC bridge forwarding defined before this L2VPN specification.

Separate eVLANs, however, may be defined with *subsets* of the eCM bridge port interfaces for independent layer 2 forwarding. In particular, customer Transparent LAN Service (TLS) is implemented by defining an eVLAN for the subscriber's L2VPN that contains only the RF Interface and the CPE Interface; a customer's TLS L2VPN is not permitted to access the eCM or any eSAFE hosts.

The eVLAN model allows a cable operator to implement Management L2VPNs for the eCM and eSAFE traffic, by defining an L2VPN with a CMIM that bridges only the RF interface, and the eCM self and/or eSAFE logical bridge interfaces.

III.1 IEEE 802.1Q and Embedded VLAN Model ¹¹²

The operation and management of a MAC layer bridge with multiple VLANs is standardized with the [802.1Q] specification. IEEE 802.1Q was first standardized in 1998 and has a standards-track MIB [RFC 4363]. The CMIM of an L2VPN can be considered to define the dot1qVlanCurrentEgressPorts bit mask of [RFC 4363]. If the eVLAN concept is adopted as the DOCSIS CM layer 2 forwarding model, [RFC 4363] already defines a rich set of objects for reporting and controlling CM layer 2 operation.

[802.1Q] was significantly upgraded in 2011. The MIB for the expanded IEEE 802.1Q has been defined in [RFC 4363].

Adopting the eVLAN forwarding model allows future DOCSIS specifications to clearly separate RF interface operation from the layer 2 filtering, forwarding, and replication to the various internal and external physical interfaces on CM-based DOCSIS devices.

[802.1Q] is extremely general purpose and sophisticated, and as a result, is also extremely complex. The specification is 327 pages long and its MIB is 99 pages. Implementing even the minimum functions specified for compliance to the IEEE 802.1Q specification, or the minimum required objects for [RFC 4364], is far more functionality and control than appropriate for the embedded MAC bridge of an L2VPN-compliant CM.

¹¹² Section revised per L2VPN-N-07.0461-2 on 7/5/07 by KN, modified per L2VPN-N-14.1186-1 on 1/16/15 by PO.

And yet, the extensive capabilities and MIBs developed by the IEEE for multiple-VLAN bridging can and should serve as a model for the future enhancement of CM layer 2 forwarding specifications. The eVLAN concept has the power to represent all of the current eDOCSIS forwarding models and can cleanly represent the L2 forwarding models for future DOCSIS specifications for IPv6 forwarding and IP multicast enhancements. IEEE802.1Q-2005, for example, defines standard management objects for performing IP protocol based VLAN classification and even individual source MAC based VLAN classification.

This specification therefore uses [802.1Q] and [RFC 4363] as informational-only conceptual guidelines for the required functionality of an L2VPN-compliant CM (and CMTS, for that matter). Future versions of this (and other) specifications may add additional layer 2 forwarding functions, and [802.1Q] and [RFC 4363] should serve as a guide for defining those functions.

As an example, the current L2VPN specification deals only with untagged packets on the eCM's logical bridge port interfaces. On the RF interface, the particular L2VPN (or as [802.1Q] terms it, the particular VLAN) for an L2PDU is always *implied* at the CMTS or CM ingress MAC Domain by the upstream service flow or downstream SAID. Future versions of this specification may introduce the concept of IEEE 802.1Q service delimiting tags on the RF Interface and/or on the CPE interface of the eCM MAC Bridge. In this case, the future specification should use concepts and MIB objects as already standardized by the industry with [802.1Q] and the IETF.

III.2 Embedded Bridge MAC Domain Service Primitives¹¹³

Because of the Service Flow capabilities of a DOCSIS RF MAC Domain, an eCM's MAC Bridge is defined to provide the following conceptual service to the RF MAC Domain:

- Downstream (RF MAC Domain to Bridge):
M_UNITDATA.request (
 L2PDU,
 eVLAN,
 user_priority)
- Upstream (Bridge to RF MAC Domain):
M_UNITDATA.indication (
 L2PDU,
 eVLAN,
 user_priority,
 ingress_port)

Where:

- L2PDU is an (untagged) Ethernet PDU with DMAC, SMAC, EtherType, and 0 to 1500 bytes of L2 payload.
- eVLAN is a local identifier for a particular eVLAN;
- user_priority is an 8-valued priority for layer 2 forwarding of the L2PDU, as defined by [802.1Q].
- ingress_port is the bridge's logical ifIndex value from which the L2PDU was received.

Downstream eCM bridge packet forwarding proceeds as follows:

1. The CM's MAC Domain (CM-MD) subcomponent receives a DOCSIS PDU from its docsCableDownstream interface that contains a non-MAC management L2PDU. The DOCSIS PDU may contain a BPI Extended Header or a Downstream Service Extended Header.
2. If the packet was encrypted with an L2VPN SAID, the CM-MD sets the requested bridge eVLAN to the one it created for the L2VPN; otherwise, the CM-MD sets the requested eVLAN to the Default eVLAN.
3. If the packet included a Downstream Service ID (DSID) that identifies the downstream service flow, the CM-MD sets the requested user_priority to the Traffic Priority parameter of the DS SF; otherwise, the CM-MD sets the requested user_priority to zero (0).

¹¹³ Various modifications throughout section per L2VPN-N-12.1065-11 on 9/19/12 by CP.

4. The CM-MD requests the MAC Bridge to forward the L2PDU on the requested eVLAN with the requested user_priority.
5. The MAC Bridge forwards the packet to an egress logical bridge port, according to the permitted egress ports as indicated in the CMIM for the eVLAN. It may flood the packet to multiple bridge ports.
6. If the MAC Bridge forwards the L2PDU to a physical CPE physical interface bridge port, it implements at least two IEEE 802.1Q Traffic Classes in order to provide QOS prioritized forwarding of layer 2 packets. CMs may implement from two to eight (8) traffic classes.
7. If the MAC Bridge forwards the L2PDU to the internal eRouter with a user_priority derived from a DS service flow Traffic Priority, the eRouter uses this value as its Traffic Importance Number of the packet.

Upstream eCM bridge packet forwarding proceeds as follows:

1. The CPE physical interface receives an L2PDU. The CPE physical interface requests the eCM MAC Bridge to forward the packet with user_priority 0 and the Default eVLAN.
2. Alternatively, an internal eSAFE device may request the MAC bridge to forward an L2PDU with an explicit user_priority and eVLAN value.
3. The MAC bridge indicates an L2PDU to be transmitted to the CM MAC Domain (CM-MD) subcomponent with an indicated eVLAN, ingress interface, and user_priority.
4. The CM-MD uses the eVLAN to select a set of Upstream Packet Classifiers; the default eVLAN will select the non-L2VPN classifiers, while any other eVLAN will select the L2VPN Upstream Packet Classifiers only for the corresponding L2VPN.
5. The CM-MD uses the indicated ingress port to match classifier rules with a CMIM criterion, and uses the indicated user_priority to match classifier rules with a User Priority Range criterion. These criteria apply to both L2VPN and non-L2VPN forwarding.
6. The CM-MD classifies the L2PDU to an upstream service flow and forwards the packet to the docsCableUpstream interface.

At this time, the L2VPN specification requires such packets to be considered to have a received user_priority of zero (0). Future versions of this specification may implement various [802.1Q] mechanisms for explicitly signaling (with priority-only tags), implicitly configuring (with default upstream user_priority), and regenerating the user-priority.

Rather than modeling an eRouter as directly outputting to the physical CPE port, the model should be modified to have the eRouter transmit to a separate eRouter CPE eVLAN that includes only the eRouter and physical CPE bridge ports. The Traffic Importance Number determined by the eRouter becomes the user_priority of the eRouter's request to transmit on the eRouter CPE eVLAN. This model makes it clear how eRouter (and any other future Layer 3 CPE forwarder) can share the physical CPE port with other forwarders to the physical CPE port in the eCM, while still maintaining QOS prioritization.

Appendix IV L2VPN Non-compliant CM Restrictions (Informative)¹¹⁴

The L2VPN service is primarily implemented at the CMTS. An operator can deploy L2VPN service using a compliant CMTS and non-compliant CMs. The restrictions when using non-compliant CMs are:

- L2VPN subscribers with non-compliant DOCSIS 1.1 and later CMs may not observe transparent forwarding of IP multicasts. This is especially troublesome when OSPF and RIPv2 advertisements are not forwarded to subscriber premise routers. Non-compliant DOCSIS 1.1 and 2.0 CMs still enforce IP Multicast forwarding rules, and so will block downstream forwarding of unjoined IP multicast groups. However, DOCSIS 2.0 CMs that implement the Static Multicast MAC parameter may be programmed to forward the desired multicast traffic. Non-compliant DOCSIS 1.0 CMs will not drop this multicast traffic. Some CM vendors also offer proprietary configurations to promiscuously forward all downstream IP multicasts.
- Unencrypted non-L2VPN layer 2 non-unicasts will leak onto L2VPN CPE networks. See Section IV.1 for a further description of this issue.
- Non-compliant CMs may not forward maximum-sized packets with a subscriber tag; i.e., of length 1522 bytes (pre3.1 DOCSIS) or 2000 bytes (DOCSIS3.1). Stacked or Tag-in-tag operation may not be possible with such CMs.
- Non-compliant CMs cannot exclude downstream L2VPN traffic from reaching the IP stacks of the embedded CMs and embedded eSAFE hosts of the L2VPN's CMs.

NOTE: *upstream* traffic from the eCMs and (usually) eSAFE hosts is blocked, preventing bi-directional unauthorized access.

- Non-compliant CMs cannot join L2VPNs dynamically, i.e., via dynamic service flow messages initiated by the CMTS after registration. Non-compliant CMs are to be statically configured to join all required L2VPNs based on the L2VPN Encodings configured in their CM configuration file or on the CMTS.¹¹⁵

IV.1 Leaking through non-compliant CMs

This specification does not specify any mechanism to prevent the leakage of non-L2VPN unencrypted traffic through *non-compliant* CMs configured for L2VPN forwarding. Table IV–1 summarizes the conditions under which downstream non-L2VPN non-unicasts can leak into a subscriber's CPE network, when a non-compliant CM is configured for L2VPN forwarding.

Table IV–1 - Non-L2VPN leaking through Non-compliant CMs configured for L2VPN

Downstream Traffic Type	DIME	
	Enabled	Disabled
Arp/DHCP Broadcasts (unencrypted)	Leaks	Leaks
Unjoined IP Multicasts, e.g., RIPv2, OSPF. (unencrypted)	DOCSIS 1.0 CM: Leaks DOCSIS 1.1 CM: Blocked	DOCSIS 1.0 CM: Leaks DOCSIS 1.1 CM: Blocked
Joined IP Multicast (encrypted when DIME enabled)	Blocked	DOCSIS 1.0 CM: Leaks DOCSIS 1.1 CM: Blocked
DSG (unencrypted always)	DOCSIS 1.0 CM: Leaks DOCSIS 1.1 CM: Blocked	DOCSIS 1.0 CM: Leaks DOCSIS 1.1 CM: Blocked

NOTE: Leaking of the unencrypted non-L2VPN broadcast traffic (ARPs and DHCP) onto an L2VPN subscriber's network is usually not a major issue for the subscriber, because such traffic is relatively low. The high-volume joined IP multicast traffic is blocked even through non-compliant CMs when it is encrypted. Even the unencrypted multicast leaking through non-compliant DOCSIS 1.0 CMs can be avoided with appropriate IP filters in the DOCSIS 1.0 CM's configuration file.

¹¹⁴ Revised per L2VPN-N-14.1135-5 on 3/31/14 by JB.

¹¹⁵ Sentence modified per L2VPN-N-12.1065-11 on 9/19/12 by CP.

Appendix V Acknowledgements (Informative)

On behalf of the cable industry and our member companies, CableLabs would like to thank the following individuals who contributed to the development of this specification:

Michael Patrick, Motorola
Harsh Parandekar, Cisco
Erich Arnold, Arris
Shengyou Zeng, Cisco
Stuart Hoggan, CableLabs
Kevin Luehrs, CableLabs
Chris Donley, CableLabs

Also, CableLabs would like to thank the numerous individuals that contributed to the development of the I10 version of this specification. In particular, we want to extend our sincere appreciation and gratitude to the following L2VPN Focus Team members:

Contributor	Company Affiliation
Mike Emmendorfer, Mircea Orban	Arris
Ed Mallette	BHN
Margo Dolas, Victor Hou, Steve Muller	Broadcom
Jason Combs	Comcast
Rajiv Asati, Prashant Pai	Cisco
Dache Beverly, Manes Patel, Todd Witten	Cox
Sasha Fertelmeister, Adam Levy	Harmonic
Greg Fisher	Hitron
Michael Cookish, Robert Faulds, Mike Patrick	Motorola
Ying Su, Victor Zuo	Shaw
Kirk Erichsen, Lance Hassan, Kevin Noll	TWC
Chris Donley, James Kim, Karthik Sundaresan	CableLabs

Appendix VI Revision History (Informative)

VI.1 Engineering Change for CM-SP-L2VPN-I02-060728

The following Engineering Change has been incorporated in CM-SP-L2VPN-I02-060728:

ECN	ECN Date	Summary
L2VPN-N-06.0277-1	6/14/06	Update L2VPN Encodings

VI.2 Engineering Change for CM-SP-L2VPN-I03-061222

The following Engineering Change has been incorporated in CM-SP-L2VPN-I03-061222:

ECN	ECN Date	Summary
L2VPN-N-06.0315-2	11/15/2006	Correct syntax and revise name for docsL2vpnPktClassL2vpnId

VI.3 Engineering Change for CM-SP-L2VPN-I04-070518

The following Engineering Change has been incorporated in CM-SP-L2VPN-I04-070518:

ECN	ECN Date	Summary
L2VPN-N-07.0417-1	4/11/2007	DUT Filtering Encoding Clarification

VI.4 Engineering Changes for CM-SP-L2VPN-I05-070803

The following Engineering Changes have been incorporated in CM-SP-L2VPN-I05-070803:

ECN	ECN Date	Summary
L2VPN-N-07.0460-1	6/13/2007	L2VPN Confirmation Codes
L2VPN-N-07.0461-2	6/13/2007	L2VPN Editorial Changes
L2VPN-N-07.0471-2	6/27/2007	DOCSIS 3.0 Compliance for L2VPN part 1

VI.5 Engineering Changes for CM-SP-L2VPN-I06-071206

The following Engineering Changes have been incorporated in CM-SP-L2VPN-I06-071206:

ECN	ECN Date	Summary
L2VPN-N-07.0504-1	7/11/2007	DOCSIS 3.0 compliance for L2VPN part 2
L2VPN-N-07.0538-2	10/24/2007	L2VPN miscellaneous corrections

VI.6 Engineering Change for CM-SP-L2VPN-I07-080215

The following Engineering Changes have been incorporated in CM-SP-L2VPN-I07-080215:

ECN	ECN Date	Summary
L2VPN-N-07.0571-1	11/28/2007	Update IMPORT clause to reference the CL-DEF-MIB
L2VPN-N-07.0592-1	1/16/2008	L2VPN Phasing Guidance

VI.7 Engineering Changes for CM-SP-L2VPN-I08-080522

The following Engineering Change has been incorporated in CM-SP-L2VPN-I08-080522:

ECN	ECN Date	Summary
L2VPN-N-08.0625-2	4/9/2008	L2VPN MAC Aging Mode

VI.8 Engineering Changes for CM-SP-L2VPN-I09-100611

The following Engineering Change has been incorporated in CM-SP-L2VPN-I09-100611:

ECN	ECN Date	Summary
L2VPN-N-10.0918-2	4/6/2010	L2VPN MPLS Update

VI.9 Engineering Changes for CM-SP-L2VPN-I10-121004

The following Engineering Change has been incorporated in CM-SP-L2VPN-I10-121004:

ECN	ECN Date	Summary
L2VPN-N-12.1065-11	8/22/2012	Updates regarding MPLS, QoS, Forwarding, S-OAM, alignment with DPoE, and misc. requested topics/features.

VI.10 Engineering Change for CM-SP-L2VPN-I11-130808

The following Engineering Change has been incorporated in CM-SP-L2VPN-I11-130808:

ECN	ECN Date	Summary	Author
L2VPN-N-13.1109-3	7/3/2013	Updates to support BGP for auto-discovery and PW signaling	Kim

VI.11 Engineering Change for CM-SP-L2VPN-I12-131120

The following Engineering Change has been incorporated in CM-SP-L2VPN-I12-131120:

ECN	ECN Date	Summary	Author
L2VPN-N-13.1120-2	8/21/2013	Removal of eSTB-DSG virtual interface from L2VPN capability	Mudugere

VI.12 Engineering Change for CM-SP-L2VPN-I13-140403

The following Engineering Change has been incorporated in CM-SP-L2VPN-I13-140403:

ECN	ECN Date	Summary	Author
L2VPN-N-14.1135-5	2/26/2014	Update to SOAM FM/PM requirements and minor corrections to other areas	Kim

VI.13 Engineering Change for CM-SP-L2VPN-I14-150305

The following Engineering Change has been incorporated in CM-SP-L2VPN-I14-150305:

ECN	ECN Date	Summary	Author
L2VPN-N-14.1186-1	9/24/2014	L2VPN Editorial Changes	Kim
