

DOCSIS® Provisioning of EPON Specifications

DPoEv1.0

DPoE MAC and Upper Layer Protocols Requirements

DPoE-SP-MULPIv1.0-I09-140807

ISSUED

Notice

This DPoE™ 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 affiliated 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:	DPoE-SP-MULPIV1.0-I09-140807			
Document Title:	DPoE MAC and Upper Layer Protocols Requirements			
Revision History:	I01 - Released 02/25/11 I02 - Released 06/07/12 I03 - Released 08/30/12 I04 - Released 11/15/12 I05 - Released 03/28/13 I06 - Released 08/08/13 I07 - Released 11/14/13 I08 - Released 03/27/14 I09 - Released 08/07/14			
Date:	August 7, 2014			
Status:	Work in Progress	Draft	Issued	Closed
Distribution Restrictions:	Author 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	DPoE Technology Introduction	9
1.2	Scope	9
1.3	Goals	10
1.4	Requirements	10
1.5	Organization of Specifications	10
1.6	DPoE Specifications	11
1.7	Reference Architecture	12
1.8	DPoE Interfaces and Reference Points	13
2	REFERENCES	15
2.1	Normative References	15
2.2	Informative References	16
2.3	Reference Acquisition	17
3	TERMS AND DEFINITIONS	18
3.1	DPoE Elements	18
3.2	Other Terms and Definitions	19
4	ABBREVIATIONS AND ACRONYMS	20
5	OVERVIEW AND THEORY OF OPERATIONS	23
5.1	MULPI Key Features	23
5.2	Technical Overview	23
5.2.1	<i>Multicast Operation</i>	23
5.2.2	<i>Network and Higher Layer Protocols</i>	23
5.2.3	<i>vCM, DPoE ONU, and CPE Provisioning and Management</i>	24
5.2.4	<i>Relationship to the Physical Plant Topology</i>	26
6	MEDIA ACCESS CONTROL SPECIFICATION	27
6.1	Introduction	27
6.1.1	<i>Overview</i>	27
6.1.2	<i>Definitions</i>	27
6.1.3	<i>Future Use</i>	28
6.2	MAC Frame Formats	28
6.2.1	<i>Generic MAC Frame Format</i>	28
6.3	MAC Management Messages	29
6.3.1	<i>DPoE OAM Messages</i>	29
7	MEDIA ACCESS CONTROL PROTOCOL OPERATION	30
7.1	Upstream Data Transmission	30
7.1.1	<i>Upstream Bandwidth Allocation</i>	30
7.1.2	<i>Upstream Transmission Request Policies and Contention Resolution</i>	30
7.1.3	<i>Upstream Service Flow Scheduling Services</i>	31
7.2	Upstream – Downstream Channel Association within a MAC Domain	31
7.3	Quality of Service	32
7.3.1	<i>Definition of a Service Flow</i>	32
7.3.2	<i>Parameters of Service Flow</i>	32
7.3.3	<i>SF-Related Requirements</i>	32
7.3.4	<i>Classifiers</i>	33
7.3.5	<i>Service Classes</i>	34
7.3.6	<i>Authorization</i>	34
7.3.7	<i>SF and Classifiers</i>	34

7.4	Data Link Encryption Support.....	35
7.4.1	MAC Messages	35
7.4.2	Framing	35
8	DATA FORWARDING	36
8.1	MEF Forwarding Requirements	36
8.1.1	Provider Bridge (PB).....	36
8.1.2	Provider Backbone Bridge (PBB).....	37
8.2	Multicast Forwarding.....	37
8.3	Requirements for IP(HSD) Forwarding	37
8.3.1	IP Serving Group (IP-SG)	38
9	DPOE SYSTEM AND DPOE ONU INTERACTION.....	39
9.1	DPoE ONU and vCM Initialization and Reinitialization.....	39
9.1.1	Scan for Downstream Channel	39
9.1.2	Continue Downstream Scanning.....	39
9.1.3	Service Group Discovery and Initial Ranging.....	39
9.1.4	Authentication.....	39
9.1.5	Establishing IP Connectivity.....	39
9.1.6	Registration with the DPoE System.....	44
9.1.7	Service IDs During vCM Initialization	45
9.2	Periodic Maintenance	45
9.3	Fault Detection and Recovery.....	46
9.3.1	MAC Layer Error-Handling	46
9.4	UNI Management	46
10	DOWNLOADING CABLE MODEM OPERATING SOFTWARE.....	48
ANNEX A	WELL-KNOWN ADDRESSES.....	49
ANNEX B	PARAMETERS AND CONSTANTS.....	50
ANNEX C	COMMON TLV ENCODINGS.....	51
C.1	[802.1ad] S-VLAN and C-VLAN Frame Classification Encodings.....	51
C.1.1	[802.1ad] S-VLAN TPID	51
C.1.2	[802.1ad] S-VLAN VID.....	51
C.1.3	[802.1ad] C-VLAN TPID	52
C.1.4	[802.1ad] C-VLAN VID.....	52
C.2	[802.1ah] I-TAG Packet Classification Encodings.....	52
C.2.1	[802.1ah] I-TAG I-TPID.....	53
C.2.2	[802.1ah] I-TAG I-SID.....	53
C.3	Top Level TLVs	53
C.4	TLV 11	55
C.5	Security (TLV 17).....	56
C.6	Classification (TLVs 22 & 23)	56
C.7	Service Flows (TLVs 24 & 25).....	58
C.8	Device Management (TLVs 38, 53 & 54)	59
C.9	TLV 43	60
C.10	[DPoE-SP-MEFv1.0] and [L2VPN] (TLVs 43.5, 45, & 65)	61
C.11	Customer (Subscriber) Management (TLV 43.7)	61
C.12	Upstream Drop Classification (TLV 60)	61
ANNEX D	ESAFE DHCP SNOOPING	63
ANNEX E	MPCP DISCOVERY PROCESSING IN DPOE NETWORKS	64
E.1	IEEE MPCP Discovery Process	64
E.2	DPoE MPCP Discovery Process.....	65
E.2.1	Discovery Using Multiple GATE messages	66

APPENDIX I	EPON MEDIA ACCESS CONTROL PROTOCOL (INFORMATIVE)	69
I.1	Timing and Synchronization	69
I.1.1	<i>MPCP Clock Synchronization</i>	69
I.1.2	<i>Loop Timing in EPON</i>	69
APPENDIX II	EPON MULTIPPOINT CONTROL PROTOCOL DATA UNITS (INFORMATIVE)	70
APPENDIX III	ACKNOWLEDGEMENTS	72
APPENDIX IV	REVISION HISTORY	73
IV.1	Engineering Changes incorporated into DPOE-SP-MULPIv1.0-I02-120607	73
IV.2	Engineering Change incorporated into DPOE-SP-MULPIv1.0-I03-120830	73
IV.3	Engineering Change incorporated into DPOE-SP-MULPIv1.0-I04-121115	73
IV.4	Engineering Change incorporated into DPOE-SP-MULPIv1.0-I05-130328	73
IV.5	Engineering Changes incorporated into DPOE-SP-MULPIv1.0-I06-130808	73
IV.6	Engineering Change incorporated into DPOE-SP-MULPIv1.0-I07-131114	73
IV.7	Engineering Changes incorporated into DPOE-SP-MULPIv1.0-I08-140327	74
IV.8	Engineering Changes incorporated into DPOE-SP-MULPIv1.0-I09-140807	74

Figures

Figure 1 - DPoE Reference Architecture	12
Figure 2 - DPoE Interfaces and Reference Points.....	13
Figure 3 - DPoE Elements	18
Figure 4 - DPoE ONU Initialization	24
Figure 5 - DPoE vCM.....	25
Figure 6 - IPv4 Only Provisioning Mode for DOCSIS-specific Registration.....	40
Figure 7 - Establishment of IPv4 Network Connectivity.....	41
Figure 8 - Registration of DPoE ONU in DPoE System	44
Figure 9 - MPCP Discovery Process Message Sequence.	64
Figure 10 - Discovery Process with RegisterStartTime, RegisterAckDelay, and RegisterDeadline.	66
Figure 11 - DPoE MPCP Discovery Process Using Multiple Auxiliary GATE Messages, REGISTER_ACK Message Arrived in Grant from Auxiliary GATE Message.	67

Tables

Table 1 - DPoE Specifications.....	11
Table 2 - DPoE Interface and Reference Point Descriptions.....	14
Table 3 - DPoE Service Flow Parameters from [MULPIv3.0] Table 7-1.....	31

This page intentionally left blank

1 INTRODUCTION

Comcast Corporation, Time Warner Cable, and Bright House Networks collaborated to develop the interoperability requirements to support business services products using Ethernet Passive Optical Network (EPON) as an access technology.

DOCSIS Provisioning of EPON (DPoE) is a joint effort of operators, vendors, and suppliers to support EPON technology using existing DOCSIS-based back office systems and processes.

Ethernet PON or EPON is an [802.3] standard for a passive optical network (PON). A PON is a specific type of multi-access optical network. A multi-access optical network is an optical fiber based network technology that permits more than two network elements to transmit and receive on the same fiber. Appendix I in [DPoE-SP-ARCHv1.0] has a more detailed explanation of multi-access optical networks.

This version of the DPoE specifications is focused on DOCSIS-based provisioning and operations of Internet Protocol (IP) using DOCSIS High Speed Data (HSD), or IP(HSD) for short, and Metro Ethernet Forum (MEF) services. DPoE Networks offer IP(HSD) services functionally equivalent to DOCSIS networks, where the DPoE System acts like a DOCSIS CMTS and the DPoE System and DPoE Optical Network Unit (ONU) to appear to act like a DOCSIS CM.

1.1 DPoE Technology Introduction¹

DPoE technology was established with the following common requirements already developed by operators. Each of the participant operators had previously selected 1G-EPON and 10G-EPON as the appropriate technology for one or more applications. EPON is a widely deployed technology with a sufficient and large supply of vendors offering a variety of products for each component of the access network. 10G-EPON technology is now becoming available and is backwards compatible with 1G-EPON. A 1G-EPON network can be incrementally upgraded to 10G-EPON, adding or replacing ONUs one at a time if required. 1G-EPON and 10G-EPON are compatible with [SCTE 174] (RFoG).

The EPON protocol [802.3] supports a centralized operator-based controller architecture, with the controller commonly referred to as the Optical Line Terminal (OLT), with low cost Layer 2 access devices (ONU). The basic service mapping architecture in EPON is to map Ethernet (or IP) frame header information (such as addresses, IP DiffServ Code Points, Ethernet Q tag, S-VLAN/C-VLAN ID, ISID, bridge address, or other marking) to a logical circuit called a Logical Link Identifier (LLID) in [802.3]. The service function is similar to that used in DOCSIS networks in many ways because it is based on a centralized scheduler and uses an LLID, which functions like a SID, supports both unicast and broadcast, and has other similarities.

At the time when development efforts around the DPoE specifications started, there were no standard management interfaces for the ongoing operations and maintenance of the network, including fault management, performance management, security, etc. Operators already had fully-working and scaled-out systems that solve these challenges for DOCSIS networks. One of the primary goals for DPoE specifications was therefore to use the existing DOCSIS back office infrastructure to scale up EPON-based business services.

1.2 Scope²

As the name suggests, the scope for this document is the MAC and upper layer protocols for DPoE Networks. The MAC in DPoE Networks is EPON. This specification does not place any additional requirements on the EPON MAC beyond the [802.3] specifications. The focus of these specifications is on the support of DOCSIS-based Operations Administration Maintenance and Provisioning (OAMP) for the MAC and upper layer protocols as specified in [MULPIv3.0].

This document identifies the requirements for all DPoE Network elements to support specific [MULPIv3.0] requirements and, only where necessary, introduces new requirements in support of the [DPoE-SP-MEFv1.0]

¹ Revised per MULPIv1.0-N- 14.0155-1 on 7/14/14 by JB.

² Revised per MULPIv1.0-N- 14.0155-1 on 7/14/14 by JB.

requirements and the [DPoE-SP-ARCHv1.0] requirements for the converged operation of IP(HSD) and Metro Ethernet services in a single DPoE Network.

1.3 Goals

Collectively, the operators started the DPoE specification development to accomplish the following objectives:

- Identify and document the common requirements for triple play services for business customers over EPON.
- Adapt DOCSIS-based back office provisioning and operations models to EPON. This is the core objective of DPoE specifications.
- Develop consensus on additional requirements above and beyond DOCSIS specifications to take advantage of the capabilities of EPON. These are focused in the area of Ethernet services and MEF integration.
- Continue to leverage the supply chain and economic benefits of a large base of suppliers and high-volume supply chain in optics, subsystems, and network systems based on a commodity EPON technology. Doing so requires adapting operator processes and networks to the EPON system rather than making any changes to the EPON systems.
- Positioning DPoE specifications to continue to leverage those same benefits for 10G-EPON.
- Work with the established EPON vendor community to assure that these strategies can be effective to mutually develop DPoE Networks, and to create a marketplace for success for multiple vendors to provide solutions for the variety of needs within the operator environment.

1.4 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.

1.5 Organization of Specifications

The DPoE specifications are organized around existing DOCSIS specifications. The purpose of matching DPoE specification documents to existing CableLabs DOCSIS, IEEE, IETF, and MEF requirements is to facilitate the mapping of services from existing DOCSIS infrastructure to DPoE infrastructure, and to provide an organization that will be easy to maintain as related (referenced) standards, recommendations, or specifications undergo independent changes.

There are two types of documents in the DPoE specifications. The first includes informative and requirements documents called specifications that detail the specific requirements for products claiming compliance with the

specifications. The DPoE specifications also include a new kind of document that does not fit into any of the above categories. The IP Network Elements [DPoE-SP-IPNEv1.0] requirements are a set of common requirements for the management of IP network elements that operators have developed, which are above and beyond the requirements in DOCSIS specifications, but are nonetheless required in DOCSIS CMTS products today. These are not specifications because no new protocols or algorithms are provided. Most of the requirements in IP NE are existing requirements based on IEEE, IETF, or other network management standards.

The DPoE documents are detailed in Section 1.6 of this document and duplicated, for reference, in each of the DPoE specifications.

1.6 DPoE Specifications³

This document is one in a series of nine (9) documents comprising the DPoE specifications. Collectively these documents represent the operators' requirements for EPON-based commercial services.

Table 1 - DPoE Specifications

Document	Document Title	Description
DPoE-SP-ARCHv1.0	DPoE Architecture Specification	DOCSIS Provisioning of EPON introduction, architecture, and narrative. Specifies fundamental architectural requirements (those that apply to more than one specification). Explains the purpose of each document below.
DPoE-SP-OAMv1.0	DPoE OAM Extensions Specification	Extensions beyond [802.3] requirements.
DPoE-SP-PHYv1.0	DPoE Physical Layer Specification	Using the EPON PHY, the DPoE PHY specification makes mandatory some options within EPON and adds some additional requirements.
DPoE-SP-SECv1.0	DPoE Security and Certificate Specification	Specifications for support for DOCSIS network and system interfaces to provide transparent support of DOCSIS device authentication, code verification, and additional security for a DPoE implementation.
DPoE-SP-IPNEv1.0	DPoE IP Network Element Requirements	Best practices and operator requirements for IP network element management and operations. This document includes CMTS-like IP router requirements. This document recommends practices not currently covered by any existing DOCSIS specifications.
DPoE-SP-MULPIv1.0	DPoE MAC and Upper Layer Protocols Requirements	Specifications for support of a subset of DOCSIS 3.0 MULPI functionality with additional EPON requirements.
DPoE-SP-MEFv1.0	DPoE Metro Ethernet Forum Specification	Specifications for Metro Ethernet services added to DOCSIS static configuration provisioning model.
DPoE-SP-OSSIV1.0	DPoE Operations and Support System Interface Specification	Specifications for support of a subset of DOCSIS 3.0 OSSI functionality with additional EPON requirements.
DPoE-SP-DEMARCv1.0	DPoE Demarcation Device Specification	Specifications to support the DEMARC Auto-Configuration process to discover and provision a demarcation device connected to a DPoE ONU.

³ Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

1.7 Reference Architecture

The DPoE reference architecture identifies the elements that a DPoE Network minimally requires to illustrate and communicate the physical hardware and logical software interfaces between the functional subsystems of the DPoE architecture. The principal elements in the architecture are the DPoE System that resides in the operator network, and the DPoE ONU which may be an off the shelf EPON ONU, EPON SFP-ONU, or an EPON ONU with additional subsystems. The remaining elements in the architecture are existing servers and systems in the operator’s network. All of the server elements have connectivity through an IP (TCP/IP) network. Transport of bearer traffic, and (in some cases) Layer 2 OAM signaling is available through either IP or Layer 2 Ethernet-based Network Interfaces.

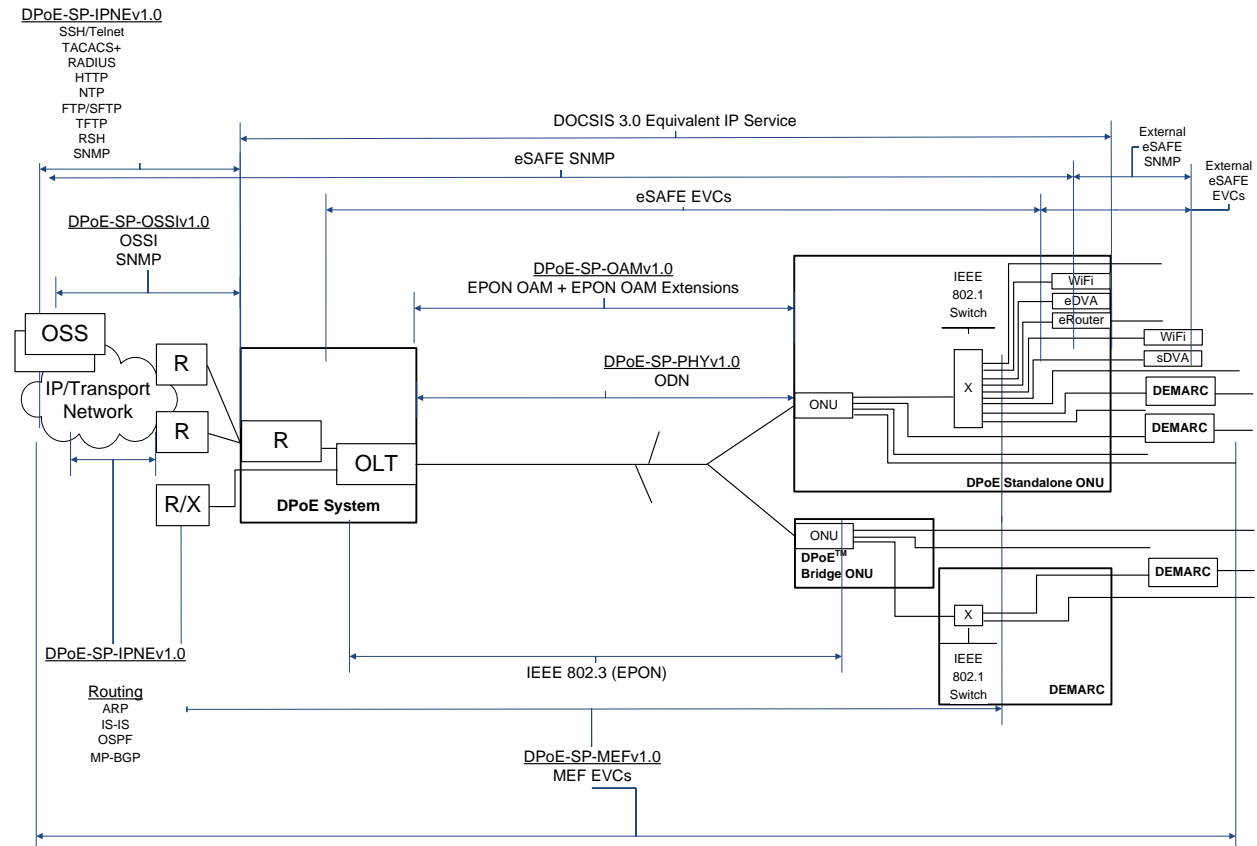


Figure 1 - DPoE Reference Architecture

1.8 DPoE Interfaces and Reference Points

The DPoE interfaces and reference points provide a basis for the description and enumeration of DPoE specifications for the DPoE architecture. Each interface or reference point indicates a point between separate sub-systems. The reference points have protocols that run across them, or have a common format of bearer traffic (with no signaling protocol). All of the interfaces are bi-directional interfaces that support two-way communications. The protocols in DPoE specifications operate within different layers based on the [802.3], [802.1], IETF, MEF, and CableLabs specifications. The C reference points are uni-directional for upstream (C_O) or downstream (C_S) classification, respectively.

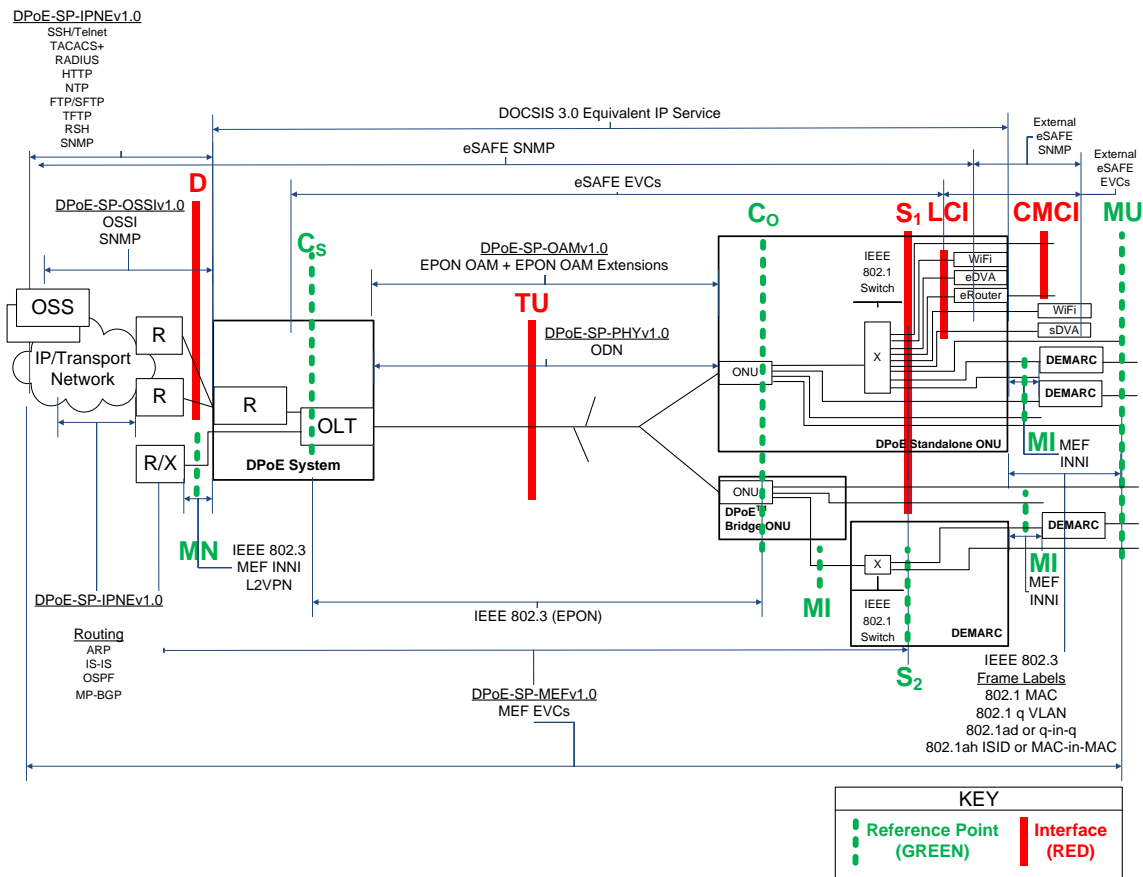


Figure 2 - DPoE Interfaces and Reference Points

Table 2 - DPoE Interface and Reference Point Descriptions

Interface or Reference Point		Interface or Reference Point Description
MN		The MN interface is an [802.3] interface for Ethernet (or MEF or L2VPN emulated) services only. It serves the role of a MEF INNI or L2VPN NSI. It is an NNI for Metro Ethernet services only.
D		The D interface is the DOCSIS IP NNI interface. It is an operator network facing interface, sometimes called a Network-to-Network Interface (NNI) or Network Systems Interface (NSI) in DOCSIS specifications. The D interface allows a DPoE System to communicate with an IP network. The D interface carries all IP management traffic including OSSI and IP NE traffic. The D interface carries all DOCSIS IP service traffic.
TU		The TU interface is a short form of expressing the interface between the DPoE System and the DPoE ONU.
C		The C reference point is used for explanation of traffic ingress to a DPoE classifier.
	C _O	The C _O reference point is used for explanation of traffic ingress to a DPoE ONU upstream classifier.
	C _S	The C _S reference point is used for explanation of traffic ingress to a DPoE System downstream classifier.
S		The S interface is an IEEE 802 interface. The S interface may be an internal interface (such as [802.3] across a GMII SERDES or XGMII interface in an SFP-ONU, SFP+ONU or XFP-ONU) or it may be an external Ethernet interface. S ₁ is an interface for a DPoE Standalone ONU. S ₂ is a reference point used for explanation of services with the DPoE Bridge ONU.
	S ₁	The S ₁ interfaces are the general case of all interfaces on a DPoE Standalone ONU. S ₁ interfaces may be CMCI, LCI, MI, or MU interfaces.
	S ₂	The S ₂ reference point is used for explanation of traffic ingress to and egress from interfaces on a DEMARC device in a DPoE System. Although there are no specifications or requirements for the S ₂ reference point, informative text refers to the S ₂ reference point to provide the full context for the use of a DPoE Bridge ONU in a DEMARC device providing Metro Ethernet services.
LCI		The Logical CPE Interface (LCI) interface is an eDOCSIS interface as defined in [eDOCSIS]. The eDOCSIS architecture is [802.1d] MAC based according to the DOCSIS 3.0 specifications; however, DOCSIS L2VPN clearly supports [802.1q] switching. In practice, therefore, the eDOCSIS interface consists of a DOCSIS classifier and [802.1] switch as illustrated. The function of a DOCSIS classifier is in part replaced by forwarding (tagging and encapsulation) in MEF, and in part covered by classifiers in this specification.
CMCI		CMCI is the DPoE interface equivalent of the DOCSIS Cable Modem CPE Interface as defined in [CMCIv3.0]. This is the service interface for DOCSIS-based IP services.
MI		MI is usually an S interface (or S reference point) that operates as a MEF INNI. A DPoE ONU that provides a MEF INNI has an MI interface. A DPoE ONU can have MU as an interface and an MI reference point on different S interfaces in a single DPoE ONU. The MI interface or reference point is an [802.3] interface (or reference point) between a DPoE ONU and a DEMARC device.
MU		MU is usually an S interface (or S reference point) that operates as a MEF UNI. A DPoE ONU that directly provides a MEF UNI (MU) interface has MU as an interface. A DPoE ONU can have MU as an interface and an MI reference point on different S interfaces in a single DPoE ONU. The MU interface or reference point is an [802.3] interface (or reference point) between a DPoE ONU or a DEMARC device and a customer's equipment.

2 REFERENCES

2.1 Normative References⁴

In order to claim compliance with this specification, it is necessary to conform to the following standards and other works as indicated, in addition to the other requirements of this specification. Notwithstanding, intellectual property rights may be required to use or implement such normative references. At the time of publication, the editions indicated were valid. All references are subject to revision, and users of this document are encouraged to investigate the possibility of applying the most recent editions of the documents listed below. References are either specific (identified by date of publication, edition number, version number, etc.) or non-specific. For a non-specific reference, the latest version applies.

- [802.1] Refers to entire suite of IEEE 802.1 standards unless otherwise specified.
- [802.1ad] IEEE Std. 802.1ad-2005TM, IEEE Standard for Local and Metropolitan Area Networks – Virtual Bridged Local Area Networks Amendment 4: Provider Bridges, May 2006.
- [802.1ah] IEEE Std. 802.1ah-2008, IEEE Standard for Local and Metropolitan Area Networks – Virtual Bridged Local Area Networks – Amendment 6: Provider Backbone Bridges, January 2008.
- [802.1d] IEEE Std 802.1dTM-2004, IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges.
- [802.1p] IEEE 802.1p (2004), LAN Layer 2 QoS/CoS Protocol For Traffic Prioritization.
- [802.1q] IEEE Std. 802.1q-2009, IEEE Standard for Local and Metropolitan Area Networks-Virtual Bridged Local Area Networks, January 2010.
- [802.3] IEEE 802.3-2012, IEEE Standard for Ethernet, December 2012
- [802.3ah] IEEE 802.3ahTM-2004: IEEE Standard for Information technology-Telecommunications and information systems-Local and metropolitan area networks-Specific requirements, Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, Amendment: Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks, now part of [802.3].
- [802.3as] IEEE Std. 802.3asTM-2006. Amendment 3 to IEEE Standard for Information technology-Telecommunications and information exchange between systems-Local and metropolitan area networks-Specific requirements-Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment 3, November 2006.
- [802.3av] IEEE 802.3avTM-2009, IEEE Standard for Information technology-Telecommunications and information systems-Local and metropolitan area networks-Specific requirements, Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment 1: Physical Layer Specifications and Management Parameters for 10Gb/s Passive Optical Networks, now part of [802.3].
- [CANN-DHCP-Reg] CableLabs' DHCP Options Registry, CL-SP-CANN-DHCP-Reg, Cable Television Laboratories, Inc.
- [CMCIv3.0] Data-Over-Cable Service Interface Specifications, Cable Modem to Customer Premise Equipment Interface Specification, CM-SP-CMCIv3.0, Cable Television Laboratories, Inc.
- [DOCSIS] Refers to entire suite of DOCSIS 3.0 specifications unless otherwise specified.

^{4 4} Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

[DPoE-SP-ARCHv1.0]	DPoE-SP-ARCHv1.0, DOCSIS Provisioning of EPON, DPoE Architecture Specification, Cable Television Laboratories, Inc.
[DPoE-SP-IPNEv1.0]	DPoE-SP-IPNEv1.0, DOCSIS Provisioning of EPON, IP Network Element Requirements, Cable Television Laboratories, Inc.
[DPoE-SP-MEFv1.0]	DPoE-SP-MEFv1.0, DOCSIS Provisioning of EPON, Metro Ethernet Forum Specification, Cable Television Laboratories, Inc.
[DPoE-SP-OAMv1.0]	DPoE-SP-OAMv1.0, DOCSIS Provisioning of EPON, OAM Extensions Specification, Cable Television Laboratories, Inc.
[DPoE-SP-OSSIV1.0]	DPoE-SP-OSSIV1.0, DOCSIS Provisioning of EPON, Operations and Support System Interface Specification, Cable Television Laboratories, Inc.
[DPoE-SP-PHYv1.0]	DPoE-SP-PHYv1.0, DOCSIS Provisioning of EPON, Physical Layer Specification, Cable Television Laboratories, Inc.
[DPoE-SP-SECv1.0]	DPoE-SP-SECv1.0, DOCSIS Provisioning of EPON, Security and Certificate Specification, Cable Television Laboratories, Inc.
[eDOCSIS]	Data-Over-Cable Service Interface Specifications, eDOCSIS Specification, CM-SP-eDOCSIS, Cable Television Laboratories, Inc.
[eRouter]	Data-Over-Cable Service Interface Specifications, eRouter Specification, CM-SP-eRouter, Cable Television Laboratories, Inc.
[L2VPN]	Data-Over-Cable Service Interface Specifications, Layer 2 Virtual Private Networks, CM-SP-L2VPN, Cable Television Laboratories, Inc.
[MULPIv3.0]	Data-Over-Cable Service Interface Specifications, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0, Cable Television Laboratories, Inc.
[RFC 2131]	IETF RFC 2131, Dynamic Host Configuration Protocol, R. Droms, March 1997.
[RFC 2132]	IETF RFC 2132, DHCP Options and BOOTP Vendor Extensions, S. Alexander, R. Droms, March 1997.
[RFC 3046]	IETF RFC 3046, DHCP Relay Agent Information Option, M. Patrick, January 2001.
[RFC 4361]	IETF RFC 4361 Node-specific Client Identifiers for Dynamic Host Configuration, T. Lemon, B. Sommerfeld, February 2006.

2.2 Informative References

This specification uses the following informative references.

[802.1ag]	IEEE Std 802.1ag–2007™, IEEE Standard for Local and metropolitan Area Networks – Virtual Bridged Local Area Networks Amendment 5: Connectivity Fault Management, December 2007.
[802.1ax]	IEEE Std. 802.1ax-2008, IEEE Standard for Local and Metropolitan Area Networks-Link Aggregation, January 2008.
[802.3ac]	IEEE Std. 802.3ac-1995™, IEEE Standard for Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications - Frame Extensions for Virtual Bridged Local Area Network (VLAN) Tagging on [802.3] Networks, January 1995. Now part of [802.3].
[DPoE-SP-DEMARCv1.0]	DPoE-SP-DEMARCv1.0, DOCSIS Provisioning of EPON, DPoE Demarcation Device Specification, Cable Television Laboratories, Inc.
[MEF 6]	Metro Ethernet Forum, MEF 6.1 Ethernet Services Definitions, Phase 2, April 2008.

[MEF 9]	Metro Ethernet Forum, Abstract Test Suite for Ethernet Services at the UNI, October 2004.
[MEF 14]	Metro Ethernet Forum, Abstract Test Suite for Traffic Management Phase 1, November 2005.
[MEF 21]	Metro Ethernet Forum, Service OAM and Requirements Framework, Phase 1, April 2007.
[MEF 26]	Metro Ethernet Forum, External Network to Network Interface (ENNI) – Phase 1, January 2010.
[OSSIV3.0]	Data-Over-Cable Service Interface Specifications, Operations Support System Interface Specification, CM-SP-OSSIV3.0, Cable Television Laboratories, Inc.
[PHYv3.0]	Data-Over-Cable Service Interface Specifications, Physical Layer Specification, CM-SP-PHYv3.0, Cable Television Laboratories, Inc.
[RFC 1918]	IETF RFC 1918, Address Allocation for Private Internets.
[RFC 2669]	DOCSIS Cable Device MIB Cable Device Management Information Base for DOCSIS compliant Cable Modems and Cable Modem Termination Systems. M. St. Johns, Ed., August 1999.
[RFC 2863]	IETF RFC 2863, The Interfaces Group MIB, K. McCloghrie, F. Kastenholz, June 2000.
[RFC 3418]	IETF RFC 3418/STD0062, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP), June 2000.
[RFC 4188]	IETF RFC 4188, K. Norseth, Ed. and E. Bell, Ed., Definitions of Managed Objects for Bridges, September 2005.
[RFC 4293]	IETF RFC 4293, Management Information Base for the Internet Protocol (IP), S. Routhier, April 2006.
[SCTE 174]	SCTE 174 2010, Radio Frequency over Glass Fiber-to-the-Home Specification.
[SECV3.0]	Data-Over-Cable Service Interface Specifications, Security Specification, CM-SP-SECV3.0, Cable Television Laboratories, Inc.
[SFF-8077i]	SFF-8077i 10 Gigabit Small Form Factor Pluggable Module, Revision 4.0, released April 13, 2004.
[SFF-8472]	SFF-8472 Specification for Diagnostic Monitoring Interface for Optical Transceivers, Revision 10.4, released January 2009.
[SFP MSA]	INF 8074i Rev 1.0, Small Form-factor Pluggable Multi-Source Agreement, released 12 May 2001.

2.3 Reference Acquisition

- Cable Television Laboratories, Inc., 858 Coal Creek Circle, Louisville, CO 80027; Phone +1-303-661-9100; Fax +1-303-661-9199; <http://www.cablelabs.com>
- Internet Engineering Task Force (IETF) Secretariat, 48377 Fremont Blvd., Suite 117, Fremont, California 94538, USA, Phone: +1-510-492-4080, Fax: +1-510-492-4001, <http://www.ietf.org>
- Institute of Electrical and Electronics Engineers (IEEE), +1 800 422 4633 (USA and Canada); <http://www.ieee.org>
- SCTE, Society of Cable Telecommunications Engineers Inc., 140 Philips Road, Exton, PA 19341 Phone: +1-800-542-5040, Fax: +1-610-363-5898, Internet: <http://www.scte.org/>
- Small Form Factor Committee (SFF), <http://www.sffcommittee.com>

3 TERMS AND DEFINITIONS

3.1 DPoE Elements

- DPoE Network** This term means the entire network described in Figure 3 from the D or MN interface to the LCI, S, MI, or MU interface (see Figure 2 for interface and reference points), depending on the service being described. In no case does the term DPoE Network ever include a DEMARC device.
- DPoE System** This term means all of the collected elements that provide the DPoE function within the operator's network facilities. This includes the EPON OLT function, DOCSIS service functions required for the D interface, Metro Ethernet service functions required for the MN interface, and IP NE element management, routing and forwarding functions specified in [DPoE-SP-IPNEv1.0]. The DPoE System is depicted in Figure 3.
- DPoE ONU** This term means a DPoE-capable ONU that complies with all of the DPoE specifications. There are two types of DPoE ONUs. These are the DPoE Standalone ONU and the DPoE Bridge ONU.
- DPoE Standalone ONU** This term means a DPoE ONU that is a standalone ONU capable of providing IP or Ethernet services directly to customer premise equipment or transport of traffic to an external DEMARC device.
- DPoE Bridge ONU** This term means a DPoE ONU that is capable of [802.1] forwarding but cannot do all of the encapsulation functions required to be a DPoE Standalone ONU. Examples include an SFP-ONU and some simple EPON chipset-only based ONUs.
- DEMARC** Short form of "Demarcation Device." This term means the device, owned and operated by the operator that provides the demarcation (sometimes called the UNI interface) to the customer. Some architectures describe this device as the CPE (as in DOCSIS, DSL, or Broadband Forum Models) or the NID (as in the MEF model).

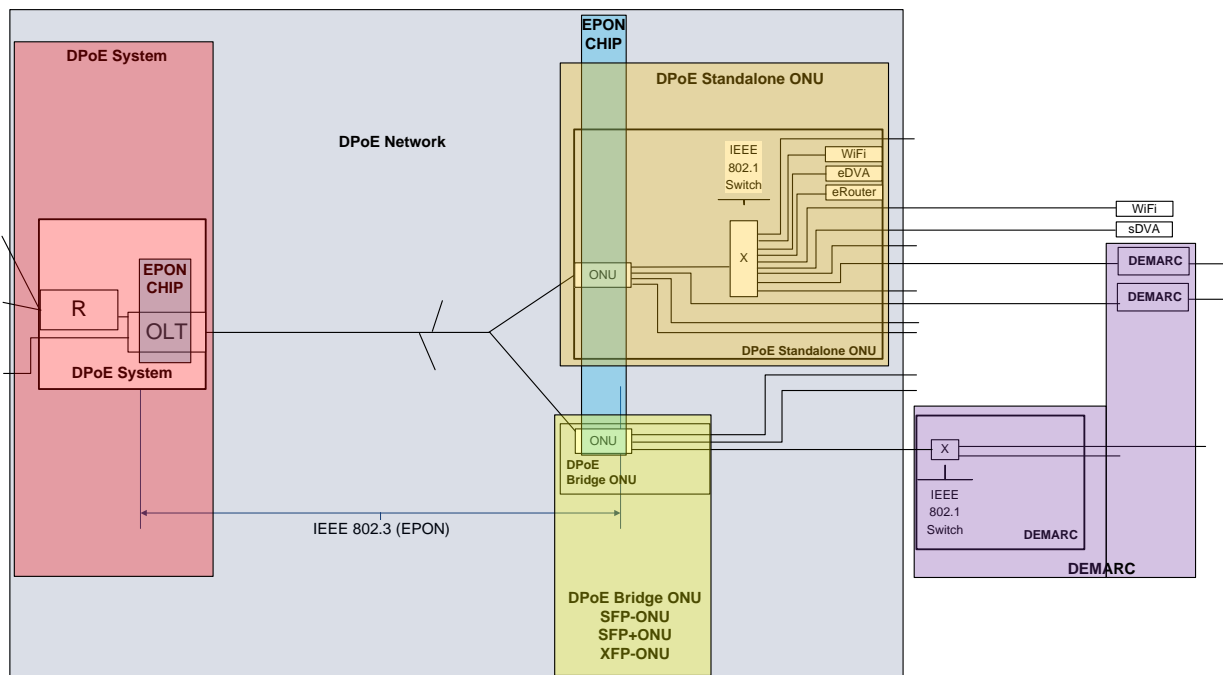


Figure 3 - DPoE Elements

3.2 Other Terms and Definitions⁵

1G-EPON	EPON as first defined in [802.3ah], now part of [802.3].
10G-EPON	EPON as defined in [802.3av], now part of [802.3].
Address Resolution Protocol	A protocol of the IETF for converting network addresses to 48-bit Ethernet addresses.
Byte	A contiguous sequence of eight bits. An octet.
Burst	A single, continuous transmission in the upstream direction originating from a single ONU, where queued customer data is transmitted towards the DPoE System at the full data rate supported by the transmission channel. Between bursts, ONUs do not transmit any data.
Cable Modem CPE Interface	CMCI as defined in [MULPIv3.0].
Classifier	A set of criteria used for packet matching according to TCP, UDP, IP, LLC, [802.1p] or [802.1q] packet fields. A classifier maps each packet to a Service Flow. A Downstream Classifier is used by the DPoE System to assign packets to downstream service flows. An Upstream Classifier is used by The DPoE ONU to assign packets to upstream service flows.
Codeword	An element of an error-correcting code used to detect and correct transmission errors.
Customer Premise Equipment (CPE)	Customer Premise Equipment as defined in [DOCSIS].
Data Link Layer	Layer 2 in the Open System Interconnection (OSI) architecture; the layer that provides services to transfer data over the transmission link between open systems (here, equal to EPON).
Data Rate	Rate Throughput, data transmitted in units of time usually in bits per second (bps). Various multipliers are used in this document, ranging from kbit/s (thousand bits per second) to Gbps (billion bits per second).
EPON Operations and Maintenance Messaging (OAM)	EPON OAM messaging as defined in [802.3] and [DPoE-SP-OAMv1.0]; Ethernet OAM is not the same as EPON OAM; Ethernet OAM is [802.1ag].
Ethernet Passive Optical Network (EPON)	Refers to both 1G-EPON and 10G-EPON collectively.
Frame	Basic data organizational unit. Here, equal to MAC frame per [802.3], Clause 4.
Logical CPE Interface	LCI as defined in [eDOCSIS].
Network Interface Device (NID)	A DEMARC device in DPoE specifications.
Upstream	The direction of transmission from the customer to the headend.

⁵ Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

4 ABBREVIATIONS AND ACRONYMS

This specification uses the following abbreviations:

BE	Best Effort
CMCI	Cable Modem CPE Interface as defined in [CMCIv3.0]
CMIM	Cable Modem Interface Mask
CoS	Class of Service
CPE	Customer Premise Equipment
CTBH	Cell Tower Backhaul
DIA	Dedicated Internet Access
DoS	Denial of Service
DPoE	DOCSIS Provisioning of EPON
DR	Default Router
DSx	Digital Signal (DS1 or DS3)
eCM	embedded Cable Modem
eDVA	embedded Digital Voice Adapter
ENNI	External Network to Network Interface
EPL	Ethernet Private Line
EPON	Ethernet Passive Optical Network; refers to both 1G-EPON and 10G-EPON collectively
EP-VLAN	Ethernet Private VLAN
eSAFE	embedded Service/Application Functional Entity
ESP	Ethernet Service Path
EVC	Ethernet Virtual Connection
E-VPL	Ethernet Virtual Private Line
EVP-LAN	Ethernet Virtual Private LAN
FEC	Forward error correction
GBd	Gigabaud
Gbps	Gigabits per second (as used in the industry)
INNI	Internal Network to Network Interface
IP	Internet Protocol
IP NE	IP Network Elements
IP(HSD)	High Speed Data Broadband Internet Access using DOCSIS
I-SID	[802.1ah] I-Component Service IDentifier
LCI	Logical CPE Interface as defined in [eDOCSIS]
LLID	Logical Link IDentifier
LTE	Logical Topology Emulation
MEF	Metro Ethernet Forum

MEN	Metro Ethernet Network
MI	MEF INNI Interface at a customer premise
MN	MEF INNI Interface to operators MEN
MPCP	Multi-Point Control Protocol
MPCPDU	MPCP Data Unit
MSC	Mobile Switching Center
MU	MEF UNI Interface
NID	Network Interface Device
NNI	Network to Network Interface
NSI	Network Systems Interface
OAM	Operations Administration and Maintenance
OAMP	Operations Administration Maintenance and Provisioning
ODN	Optical Distribution Network
OLT	Optical Line Termination
ONU	Optical Network Unit
OSC	Optical Splitter Combiner
P2MP	Point to Multi-Point
P2P	Point-to-Point
P2PE	Point-to-Point Emulation
PB	Provider Bridging [802.1ad]
PBB	Provider Backbone Bridging [802.1ah]
PCS	Physical Coding Sublayer
PHY	Physical Layer
PMA	Physical Medium Attachment
PMD	Physical Media Dependent (Sublayer)
PON	Passive Optical Network
QoS	Quality of Service
R	IP Router
RS	Reconciliation Sublayer
RTPS	Real Time Polling Service
SCB	Single Copy Broadcast
sDVA	Standalone Digital Voice Adapter
SFP	Small Form-factor Pluggable
SFP+	Small Form-factor Pluggable Plus (+)
TPID	Tag Protocol Identifier
UNI	User Network Interface
vCM	Virtual Cable Modem

WSC	Wireless Switching Center
X	IEEE Ethernet Switch (Generic)
XFP	X Form-factor Pluggable

5 OVERVIEW AND THEORY OF OPERATIONS

5.1 MULPI Key Features⁶

DPoE specifications introduce a number of features that build upon features defined in the DOCSIS 3.0 specifications, as well as EPON specifications [802.3], together with the series of other relevant specifications as listed in Section 2. This specification includes the following key new features for the MAC and Upper Layer Protocols Interface as compared to the DOCSIS 3.0 version.

- Downstream Channel transmission is operated using Time Division Multiplexing (TDM) transmission over Ethernet Passive Optical Networks (EPONs) with Intensity Modulated (IM) signal. Channel Bonding is not supported in DPoE Networks, and bandwidth is assigned to individual links / circuits on demand via the Dynamic Bandwidth Allocation (DBA) operating in the DPoE System. DPoE specifications, building on 1G-EPON [802.3], offer an effective bandwidth of approximately 960 Mbit/s (already accounting for transmission overhead, excluding optional Forward error correction (FEC) parity), while DPoE specifications, building on 10G-EPON [802.3], offer an effective bandwidth of approximately 8.9 Gbps (already accounting for transmission overhead, including mandatory FEC parity etc.). Downstream channel in DPoE Networks provide broadcast capability inherent for Point to Multi Point (P2MP) passive architecture of EPON.
- Upstream Channel transmission is operated using Time Division Multiple Access (TDMA) transmission over EPONs with IM signal, where several DPoE ONUs connected to a single DPoE System time-share a single receiver to the upstream medium. Channel Bonding is not supported in DPoE Networks, and bandwidth is assigned to individual DPoE ONUs / logical entities on demand via the DBA operating at the DPoE System. DPoE specifications, building on 1G-EPON [802.3], offer an effective bandwidth of approximately 920 Mbit/s (already accounting for transmission overhead, including typical band gaps, excluding optional FEC parity), while DPoE, building on 10G-EPON [802.3], offers an effective bandwidth of approximately 8.6 Gbps (already accounting for transmission overhead, including mandatory FEC parity and band gaps, etc.).

5.2 Technical Overview

This specification defines the MAC layer protocols of the DPoE System, as well as requirements for upper layer protocols (e.g., IP, DHCP, etc.) operating on top of the EPON MAC. DPoE specifications introduce new MAC layer features beyond what are included in earlier versions of DOCSIS specifications, inherently reusing EPON MAC definitions for transmission of Ethernet encapsulated data over P2MP passive optical links.

EPON, and therefore DPoE specifications, do not support DOCSIS MAC-specific functions, such as:

- DOCSIS Dynamic Quality of Service (QoS) establishment and two-phase activation process;
- DOCSIS-specific load balancing;
- DOCSIS channel bonding in upstream and downstream channels;
- frame fragmentation at the transport layer;
- prioritized data forwarding through the CM;
- D3.0 Multicast QoS and Group Service Flows (GSF).

5.2.1 Multicast Operation

This functionality is not included in the scope of this version of the DPoE specifications.

5.2.2 Network and Higher Layer Protocols

The DPoE System MUST perform (Ethernet) MAC Layer bridging and Network Layer routing of data traffic. The DPoE ONU MUST only perform MAC layer bridging of data traffic. However, both DPoE System and DPoE ONU are network-layer and transport-layer aware. Specifically, the DPoE System and DPoE ONU support classifying

⁶ Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

user traffic, based on network layer and transport layer criteria, for purposes of providing Quality of Service and packet filtering.

Additionally, the DPoE System MUST support the following protocols for operation and management:

- SNMP (Simple Network Management Protocol)
- TFTP (Trivial File Transfer Protocol), which is used by the DPoE System for downloading operational software and configuration information
- DHCPv4 (Dynamic Host Configuration Protocol) is used by the DPoE System to obtain IP addresses and other configuration for the purposes of DPoE ONU management proxy.

5.2.3 vCM, DPoE ONU, and CPE Provisioning and Management

5.2.3.1 Initialization, Provisioning and Management of CMs⁷

During initialization, the DPoE ONU goes through a number of steps before becoming fully operational on the DPoE Network. The full initialization comprises the same four fundamental stages specified for a DOCSIS CM:

1. Topology resolution and physical layer initialization
2. Authentication and encryption initialization
3. IP initialization
4. Registration (MAC layer initialization)

The DPoE ONU initialization is shown in Figure 4.




Figure 4 - DPoE ONU Initialization

⁷ Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

The first stage, topology resolution and physical layer initialization, is specified in the [DPoE-SP-PHYv1.0] specification. The [DPoE-SP-PHYv1.0] specification describes the DPoE relationship to EPON system specifications.

The second stage, authentication and encryption, is specified in the [DPoE-SP-SECv1.0] specification. The [DPoE-SP-SECv1.0] specification describes how security is implemented in a DPoE Network.

The third stage, IP Initialization, requires the assignment of an IPv4 address enabling DPoE System management of the DPoE ONU. Since the DPoE ONU does not contain an IP stack and is not directly addressable using IP, this specification requires that the DPoE System MUST obtain an IP address and modem provisioning file from the OSS provisioning systems, on behalf of the DPoE ONU, as part of the modem registration process.

As described in the [DPoE-SP-OSSv1.0] specification, the DPoE System MUST provide management capabilities on behalf of the DPoE ONU for all IP-based management functions when the OSS management systems direct management requests to a given DPoE ONU. The concept of a virtual Cable Modem (vCM) is used in this specification to represent the IP-addressable management entity maintained and controlled within the DPoE System. See Figure 5.

When the DPoE System receives management requests for a vCM, it converts those requests into the appropriate DPoE OAM requests (if the given management request requires an eOAM exchange between the DPoE System and the DPoE ONU), and sends the Operations Administration and Maintenance (OAM) requests to the DPoE ONU as needed. See [DPoE-SP-OAMv1.0] for a full description of the DPoE OAM messaging. The vCM is used to map requirements that were previously required of the DOCSIS Cable Modem to requirements on the DPoE System.

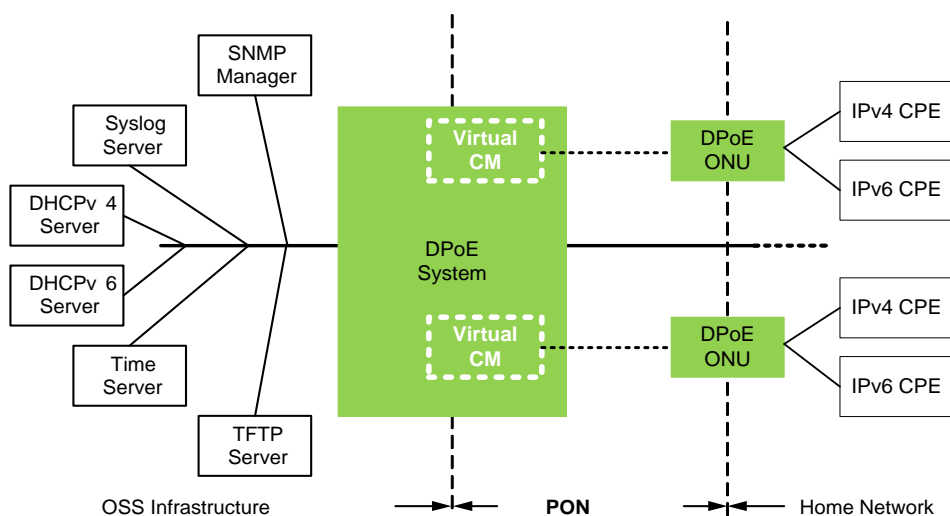


Figure 5 - DPoE vCM

The DPoE System uses DHCPv4 to acquire an IPv4 address, followed by TFTP to obtain DPoE ONU operational parameters. To facilitate compatibility with existing provisioning systems, this process is identical to the DOCSIS CM provisioning process and is further described in Section 9 of this document.

The fourth stage, registration, involves the DPoE System processing the DOCSIS configuration file. The DPoE System validates the contents and configures the DPoE System and DPoE ONU based on the service provisioning information that it received. The vCM is used by the DPoE System to store the registration state, as well as the configuration of the DPoE ONU.

After the DPoE System completes initialization, the vCM is a manageable network element in the operator's IP network. The vCM supports SNMP (as mentioned above), and responds to queries directed to the IPv4 address that it acquired during initialization.

5.2.3.2 Initialization, Provisioning, and Management of CPEs⁸

DOCSIS specifications assume the use of DHCP for provisioning of CPE devices (per [DPoE-SP-ARCHv1.0]). To that end, the DPoE System MUST support a DHCP Relay Agent that allows the operator to associate a CPE IP Address request with the customer DPoE ONU MAC Address. This feature is also used as the basis of a mechanism that prevents spoofing of IP Addresses.

If a CPE client DCHPDISCOVER is received by the DPoE System on an upstream service flow that is configured in an IP-SG, the DPoE System relay agent MUST set the giaddr field in accordance with the IP parameter specified in the IP-SG.

The DPoE System supports the provisioning of CPE devices with an IPv4 address.

5.2.4 Relationship to the Physical Plant Topology

DPoE Networks use an all fiber passive Optical Distribution Network (ODN). There are no active elements in the ODN. The only elements in a DPoE Network ODN are optical splitter/combiners, connectors, and fiber.

In the DPoE specifications, the concept of the DPoE MAC Domain is introduced as a substitute for the DOCSIS MAC Domain. The DPoE MAC Domain is used for the exclusive purpose of compatibility with [OSSIPv3.0] functions.

The DPoE specifications implement the [DOCSIS] protocols primarily through the proxy of functions on the DPoE System, which operates a vCM in place of a CM to emulate the function of the CM for management purposes in a DPoE Network. The DOCSIS MAC domain is used to directly manage CMs in the MAC domain, without respect to the IP address of the CM. Likewise, the DPoE Network directly manages the flow of services at Layer 2 using the EPON LLID (in place of the DOCSIS SID). A DPoE System MAC Domain MUST include all of the DPoE ONUs on a TU interface.

⁸ Revised per MULPIv1.0-N-14.0117-2 on 2/19/14 by JB.

6 MEDIA ACCESS CONTROL SPECIFICATION

6.1 Introduction

6.1.1 Overview

DPoE Networks use the EPON MAC. The additional requirements beyond the EPON specifications are provided in this section and include the following functions of the MAC, MAC Clients, and MAC Control sublayers:

- Dynamic Bandwidth Allocation (DBA) controlled by the DPoE System for the upstream channel.
- Downstream channel operated using TDM technique.
- Upstream channel operated using TDMA technique with centrally controlled medium access.
- High bandwidth efficiency through native support for Ethernet encapsulation and limitation of protocol conversions between the CPE and Ethernet-based transport network.
- Native support for variable packet lengths (see [DPoE-SP-ARCHv1.0] for more details on the size of supported frames). Note that, in either case, the maximum payload size of the Ethernet frame remains exactly the same and is equal to 1500 bytes. This feature combines naturally with variable-size packets produced by TCP/IP carrying the majority of today's network traffic.
- Strong QoS support through
 - support for Bandwidth and Latency Guarantees,
 - packet classification,
 - dynamic, on-demand bandwidth allocation for increased on-demand throughput,
 - inherent support for various levels of fairness for bandwidth sharing.
- Strong link layer security as described in [DPoE-SP-SECv1.0].

6.1.2 Definitions

6.1.2.1 MAC Frame⁹

This is an Ethernet frame as defined in Clause 4 [802.3] with extensions introduced in Clause 65 [802.3] and Clause 76 [802.3] to support Point-to-Point (P2P) emulation on P2MP physical medium. DPoE Networks do not modify the Ethernet frame format in any way.

6.1.2.2 Time Division Multiplexing

This is a mode of transmission between the DPoE System and DPoE ONUs (in downstream) in which the data channel is divided into a number of smaller transmission slots (bits), any number of which can be assigned to any of the receiving stations depending on the data pending transmission at the DPoE System. Data in DPoE Networks is sent in MAC Frames, comprising an organized structure of data as defined in Section 6.1.2.1.

In the downstream channel, there is only one transmitter active at any time (DPoE System), and multiple receivers (DPoE ONUs), hence no collision avoidance mechanism is needed.

6.1.2.3 Time Division Multiple Access

This is a mode of transmission between the DPoE ONUs and the DPoE System (in upstream), in which individual stations take possession of the data channel for a strictly controlled period of time and send data towards the DPoE System using TDM transmission. Data in DPoE Networks is sent in MAC Frames, comprising an organized structure of data as defined in Section 6.1.2.1.

In the upstream channel, there are multiple transmitters that can be active at any time, and only one receiver; hence, a collision avoidance mechanism is needed, in which a central medium access controller (DPoE System) assigns

⁹ Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

transmission opportunities to individual DPoE ONUs in a pre-defined and non-colliding manner. Bandwidth assignment is dynamic and depends on the SLAs, supported services, fairness, etc., subject to definition by the service provider.

6.1.2.4 LLID¹⁰

In DPoE specifications, as in the EPON standard, an LLID is used to identify a logical link between the DPoE System and a DPoE ONU. The LLID is a 15-bit field transmitted in the preamble of both upstream and downstream frames. During the registration process, the DPoE System assigns one or more LLIDs to a DPoE ONU. In the downstream direction, if the frame received by a DPoE ONU contains an LLID that matches an LLID previously assigned to a DPoE ONU, then the DPoE ONU accepts the frame and forwards it to its MAC layer. Conversely, if the frame received by a DPoE ONU contains an LLID that does not match an LLID previously assigned to a DPoE ONU, the frame is rejected by the DPoE ONU. In the upstream direction, based on the LLID in an incoming frame, the DPoE System forwards frames to the appropriate MAC entity. In this way, a point-to-point emulation over the point-to-multipoint PON architecture is created.

The LLID is also used to represent a traffic-bearing entity that requires upstream bandwidth allocation by the DPoE System. Through a queue occupancy reporting mechanism defined in the EPON standard, the DPoE ONU informs the DPoE System how much data is waiting to be transmitted for a particular LLID. The DPoE System, in turn, schedules the LLID for upstream transmission.

In 1G-EPON, one LLID (value 0x7FFF) is reserved for broadcast transmission (flooding data to all DPoE ONUs) in the downstream channel and transmission of unregistered DPoE ONUs in the upstream channel. In 10G-EPON, a different LLID is reserved (0x7FFE) for the same purpose to facilitate coexistence between both systems. Additionally, a range of LLIDs is left reserved for future use - 0x7FFD–0x7F00 (see Table 76-4 in [802.3] for details).

6.1.3 Future Use

There are no fields in the EPON MAC frame marked as "for future use".

Ethernet frame format is always extended in a backward compatible way, thus guaranteeing that existing deployed equipment can remain standard compliant. An example of such an extension is the [802.3as], which introduced support for larger Ethernet frames of up to 2000 bytes (so-called envelope frames) while maintaining the currently existing frame size limitations for deployed equipment.

6.2 MAC Frame Formats¹¹

6.2.1 Generic MAC Frame Format

In the scope of this document, the term Generic MAC Frame is considered to be equal to the term Ethernet MAC frame as defined in Section 6.1.2.1. The only normative document defining the Ethernet MAC frame format in use in the DPoE System is [802.3], Clause 4.

The DPoE System **MUST** support the maximum Ethernet frame size of at least 1600 bytes for 1G-EPON. The DPoE System **MUST** support the maximum Ethernet frame size of at least 2000 bytes [802.3as] for 10G-EPON. The DPoE System **SHOULD** support the maximum Ethernet frame size of at least 4000 bytes for 1G-EPON. The DPoE System **SHOULD** support the maximum Ethernet frame size of at least 9000 bytes for 10G-EPON. The D-ONU **MUST** support the maximum Ethernet frame size of at least 1600 bytes for 1G-EPON. The D-ONU **MUST** support the maximum Ethernet frame size of at least 2000 bytes for 10G-EPON. The D-ONU **SHOULD** support the maximum Ethernet frame size of at least 4000 bytes for 1G-EPON. The D-ONU **SHOULD** support the maximum Ethernet frame size of at least 9000 bytes for 10G-EPON. These maximum sizes above indicate the size for the whole Ethernet frame including the preamble, DA, SA, Type/Size, Payload, FCS and any [802.1q] tags (encapsulation) there may be.

¹⁰ Revised per MULPIv1.0-N-12.0034-1 on 5/29/12 by JB.

¹¹ Revised per MULPIv1.0-N-14.0134-1 on 5/27/14 by JB

6.3 MAC Management Messages

Messages defined in [MULPIv3.0], subsections 6.4.1 through 6.4.36, are not supported by the DPoE System.

6.3.1 DPoE OAM Messages¹²

Operation of the [802.3], Clause 57 OAM protocol on a generic Ethernet interface (regardless of whether it is a legacy P2P full-duplex or EPON P2MP link) does not affect standard data transmission, having a minimum impact on bandwidth available for services. The OAM protocol relies on a "slow protocol" with very limited bandwidth consumption, generating at most 10 frames per second, and by definition, it is not required for normal link operation, but rather for its maintenance and fault detection. The OAM protocol can be implemented in hardware or software, thus providing the desired media independence and flexibility required, especially for legacy equipment where hardware changes are highly unwelcome and software alternations are limited in scope. OAM frames target the slow protocol MAC address (standard defined) and are intercepted by the MAC sublayer, and thus do not propagate across multiple hops in an Ethernet network, assuring the OAM protocol data units (OAMPDUs) affect only the operation of the OAM protocol itself, while leaving the contents of the customer frames unaltered.

The main supported OAM features and functionalities include:

- Discovery process – The OAM Discovery Process is the first phase of [802.3] OAM protocol, and its basic functionality is limited to identifying the individual devices in the given network domain as well as their OAM capabilities. In [802.3], upon powering up, a device enters a discovery state and attempts to send the Information OAMPDU to its link peer, thus establishing the local link information path, which will be used further on for exchange of more specialized OAM frames. In the case of standard Ethernet OAM, the discovery process relies on the Information OAMPDUs, which are propagated in the given network and trigger all connected OAM-enabled devices to issue their OAM capabilities information, which will be encapsulated in the appropriate frames and delivered to other link peer stations.
- Link performance monitoring – The OAM link monitoring administration tools target detection and identification of link faults, where the detection mechanism utilizes the Event Notification OAMPDU, sending link state-related events to a Link Partner OAM entity, thus relaying the information on the potential link problems. If the link partner happens to be SNMP-enabled, an SNMP trap could pass the OAMPDU to a remote entity. There are a number of standard defined error events, helping to diagnose the current state of the link.
- Remote fault detection – Describing means of detecting and handling compromised links in any underlying Ethernet network infrastructure
- Remote loopback – In the loopback mode, every Ethernet frame received is transmitted back on that same port except for OAMPDUs and pause frames, which provide flow control and OAM functionality. This particular functionality helps network administrators ensure and measure link quality during installation or troubleshooting stages, when no standard data exchange occurs and the given links are subject to testing and quality evaluation. The remote loopback session requires a periodic exchange of OAMPDUs messages; otherwise, the OAM session is interrupted and all link peer stations transit into the standard transmission mode. It is interesting to note that any OAM-enabled station with a link in active mode (as opposed to passive mode) can force its link peer station into the remote loopback mode simply by sending a loopback control OAMPDU. The loopback command is acknowledged by responding with an Information OAMPDU with the loopback state indicated in the state field.
- MIB variable retrieval – Providing management information look-up from a remote database, delivering required OAM-specific information on the given network structure
- Organization-specific enhancements – System vendors have chosen to utilize organization-specific extensions to the standard Ethernet OAM protocol to implement additional and extended events, include additional information during the discovery phase, or even develop a completely proprietary OAM protocol, while maintaining the general framework compatibility with the standard IEEE-compliant OAM.

Detailed specification of both [802.3], Clause 57 compliant OAM, as well as DPoE-specific extended OAM (eOAM), can be found in [DPoE-SP-OAMv1.0] document, being part of the DPoE specifications.

¹² Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

7 MEDIA ACCESS CONTROL PROTOCOL OPERATION¹³

The DPoE specifications rely on the respective [802.3] specifications and do not introduce any changes to their stipulations. Bandwidth allocation in DPoE Networks is managed through the allocation of LLIDs as equivalent to [DOCSIS] SFIDs. Bandwidth allocation for each LLID is controlled using Multi-Point Control Protocol (MPCP) as described herein.

The DPoE specifications provide additional requirements, included in the following subsections, which are not covered by the respective [802.3] specifications, providing definitions of functions outside the scope of [802.3].

7.1 Upstream Data Transmission

In EPON, in the upstream direction (from users to network), due to the directional properties of passive splitter/combiner devices deployed in the ODN, data packets transmitted from any of the connected DPoE ONU will only reach the DPoE System and not other DPoE ONUs. In this way, the properties of the upstream EPON channel are similar to those of a standard P2P link.

However, unlike a true P2P architecture, in the DPoE Network all DPoE ONUs belong to a single collision domain. This means that data packets from different DPoE ONUs transmitted simultaneously still may collide, since they are contending for access to a single DPoE System TU interface. Therefore, in the upstream direction, the DPoE Network needs to employ some form of the medium access arbitration mechanism to avoid data collisions, and fairly share the channel capacity among DPoE ONUs, subject to specific user service agreements, fairness rules imposed in the network, etc.

7.1.1 Upstream Bandwidth Allocation

Definition of the Dynamic Bandwidth Allocation (DBA) mechanism is outside of the scope of the DPoE specifications. Upstream Bandwidth Allocation mechanism in DPoE System relies on the MPCP functionalities (i.e., GATE MPCPDU and REPORT MPCPDU), as described in more detail in Appendix II.

7.1.2 Upstream Transmission Request Policies and Contention Resolution

In both DOCSIS and EPON MAC layers, the upstream direction transmission does not start until a DPoE ONU gets a grant/gate for transmission from DPoE System.

There are two upstream granting policies defined in DOCSIS/EPON MAC layer. One is "request-based" – when DPoE System gives grants as a consequence of data transmission requests made from DPoE ONU. Another is "unsolicited mode" – when the DPoE System gives grants based on its own policies without consideration for requests from DPoE ONU, usually on a periodic basis. In DOCSIS specifications, such a mechanism is called Unsolicited Grant Service (UGS), while in EPON this mechanism is called Time Division Multiplexing (TDM). This version of the DPoE specification does not require support of UGS services.

When using a "request-based" granting policy, a DPoE System will not give grants until a bandwidth request (REPORT MPCPDU) is received from DPoE ONU for transmission of data. Therefore, for proper support of QoS, it is very important to give DPoE ONU the opportunity to transmit requests on a timely basis. The DPoE System guarantees timely reception of the bandwidth requests from the connected DPoE ONUs via transmission of GATE MPCPDUs with the report request bit set.

DOCSIS networks may use contention-based request scheduling policies, where requests from different CMs may be transmitted during the same time, and therefore, may collide and need to be retransmitted at a later time. A "contention resolution" mechanism is needed to guarantee that individual CMs can deliver data to the CMTS. Obviously, such a mechanism is statistical, introduces additional transmission delay, and therefore, is not well suited for real-time services.

Contention-based request scheduling is implemented in DOCSIS networks to be able to support large numbers of CMs in the same MAC domain with limited bandwidth overhead. The DPoE Network, on the other hand, does not

¹³ Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

use contention-based request scheduling in the MAC layer. Every request transmission in the DPoE Network is always contention-less and, therefore, is delay bound.

7.1.3 Upstream Service Flow Scheduling Services

[MULPIv3.0] defines five Service Flow Scheduling Types: BE, RTP, NRTP, UGS, and UGSAD. All types could be emulated using the scheduling mechanisms inherent in EPON; however, this version of the DPoE specifications only mandates the support of two types.

A DPoE System MUST support the Service Flow Scheduling Type (TLV 24.15) values indicating Real Time Polling Service (RTPS) and Best Effort Service (BE). A DPoE System MAY support other Service Flow Scheduling Type values in a vendor-specific manner. If a DPoE System does not support a specific Scheduling Type, and it sees the corresponding value in a CM configuration file, it MUST disallow registration of the vCM.

The following table details which parameters are applicable for an Upstream Service Flow, according to its configured Upstream Scheduling Service Type.

Table 3 - DPoE Service Flow Parameters from [MULPIv3.0] Table 7-1¹⁴

Service Flow Parameter as defined in DOCSIS MULPI	Best Effort	Real-Time Polling
Miscellaneous		
Traffic Priority	Optional Default = 0	N/A ¹
Upstream Scheduling Service Type	Optional Default = 2	Mandatory
Request/Transmission Policy	Optional Default = 0	Optional Default = 0
Maximum Rate		
Max Sustained Traffic Rate	Optional Default = 0	Optional Default = 0
Max Traffic Burst	Optional Default = 12800	Optional Default = 12800
Minimum Rate		
Min Reserved Traffic Rate	Optional Default = 0	Optional Default = 0
Polls		
Nominal Polling Interval	N/A ¹	Mandatory
Table Notes:		
¹ N/A means not applicable to this service flow scheduling type.		

A DPoE System MUST deny any request for a service flow containing a parameter that is not applicable for the given service flow scheduling type. Note that a DPoE System MUST support a Nominal Polling Interval (TLV 24.17) parameter in a CM configuration file where Real Time Polling Service is specified. This parameter defines the minimum time between poll requests. The DPoE System is not required to set up an EPON polling interval exactly as configured, and it may use a value that approximates the configured one. However, the DPoE System MUST reflect the actual value (approximated by the system) in the docsQos3ParamSetNomPollInterval MIB variable. How a DPoE System chooses an EPON polling interval for a Best Effort Service flow is vendor-specific.

7.2 Upstream – Downstream Channel Association within a MAC Domain

As described in Section 5.2.4, a single TU interface (which contains a single upstream and downstream channel) defines a DPoE MAC Domain.

¹⁴ Revised per MULPIv1.0-N-12.0012-1 on 3/27/12 by JB.

7.3 Quality of Service

7.3.1 Definition of a Service Flow

The DPoE Network supports QoS through the concept of a Service Flow (SF), which is defined as a DPoE MAC-layer transport service that provides unidirectional transport of frames, transmitted in the upstream direction by the DPoE ONU, or in the downstream direction by the DPoE System. An SF is characterized by a set of QoS Parameters provisioned via the CM Configuration file, a Service Class definition in the DPoE System, or some combination of both.

The Quality of Service Parameter Set Type (TLV 24/25.6) included in the CM configuration file defines the SF states for which the QoS Parameters apply. The DPoE System MUST only support a value of "7 – bitmask having 1 for all Provisioned, Admitted and Active states", which means that the configured QoS parameters apply to all these three states. All other QoS Parameter Set values MUST cause registration to fail at the DPoE System.

7.3.2 Parameters of Service Flow¹⁵

A Service Flow is characterized by the following set of parameters. Unless stated otherwise, these parameters are to be interpreted as defined by the [MULPIv3.0] specification.

- Service Flow ID (SFID) serves as the principal identifier of the Service Flow in the DPoE System and vCM.
- The DPoE System MUST support the following QoS Parameters for DPoE Service Flows:
 - The Maximum Sustained Traffic Rate (TLV 24/25.8) parameter.
 - The Maximum Traffic Burst (TLV 24/25.9) parameter. This parameter has a minimum value of 1600 Bytes and a default value of 12800 bytes.
 - The Minimum Reserved Traffic Rate (TLV 24/25.10) parameter.
 - The DPoE System MUST support Data Rate Unit Setting (TLV 24/25.41) Parameters as defined in [MULPIv3.0]
 - The Request Transmission Policy (TLV 24.16) parameter to control EPON scheduling behavior for Upstream Service flows. Only Bit #4, the "No Piggyback Bit", is to be supported as a way to disable MPCP Force Report behavior, where REPORTs are piggybacked with data frames.

The DPoE System MUST ignore the Maximum Concatenated Burst parameter for DPoE Service Flows.

In addition to the parameters mentioned above, for IP(HSD) services the DPoE System MUST support the following parameters:

- A Service Flow Scheduling Type (TLV 24.15) value of "Best Effort" for all upstream service flows.
- The Traffic Priority (TLV 24/25.7) parameter.

In addition to the parameter mentioned above, for Metro Ethernet services the DPoE System MUST support the following parameter :

- A Service Flow Scheduling Type (TLV 24.15) value of "Real-Time Polling" for all upstream service flows.
- The Nominal Polling Interval (TLV 24.17) for "Real-Time Polling" upstream service flows.

For detailed information on mapping MEF QoS parameters into DOCSIS QoS parameters, see [DPoE-SP-MEFv1.0].

7.3.3 SF-Related Requirements

An SF within the DPoE System MUST be mapped into one of the available LLIDs, where the mapping between the SF and the LLID is a vendor-specific process and outside the scope of this specification.

¹⁵ Revised per MULPIv1.0-N-12.0012-1 on 3/27/12 and per MULPIv1.0-N-14.0127-1 on 2/27/14 by JB.

An SF begins its existence when the DPoE System assigns an SFID to it. The SFID serves as the principal identifier for the SF within the DPoE ONU and DPoE System. Within a DPoE System, any SF MUST have an assigned SFID, together with the set of QoS Parameters, as defined above. An SF may not exist without an assigned SFID.

The DPoE System MUST support parsing and processing of CM configuration files.

The DPoE System MUST support assignment of SIDs to upstream service flows. The SID is an important management identifier in DOCSIS networks. SID Clusters defined in [MULPIv3.0] are not supported in this specification.

In [MULPIv3.0], an SF can be in one of several states – provisioned, admitted, and active. For each of those states, the flow may have a different set of values of QoS parameters. However, the DPoE System MUST only support SFs that are provisioned, admitted, and active (Quality of Service Parameter Set Type (TLV 24/25.6) value of 7). There are no provisions for deferred admission and/or activation of SFs. Admission and activation of SFs in the DPoE System MUST take place when the given SF is provisioned.

7.3.4 Classifiers¹⁶

In the DPoE Network, a Classifier is defined as a set of matching criteria applied to each frame entering the DPoE Network, composed of frame matching criteria (see below for details) and a Classifier priority. Matching criteria not specified in a classifier will not be given consideration while classifying frames. If no Classifier Rule Priority is specified, the DPoE System MUST use the default Rule Priority of 0. The DPoE System MUST support assignment of classifier Ids (CID) to individual classifiers. The CID is an important management identifier as specified in [MULPIv3.0].

Additionally, a Classifier includes a reference to the target SF, to which the given selected frames are attributed if a frame meets the specified Classifier criteria.

The DPoE System MUST support classification of frames into downstream Service Flows. The DPoE ONU MUST support classification of frames into upstream Service Flows. The classification behavior follows the requirements specified in [MULPIv3.0]. The DPoE Network MUST NOT support upstream classification policing at the DPoE System.

The DPoE ONU MUST support the Upstream Drop Classifiers (TLV 60).

The DPoE ONU supports the following combinations of classifiers:

- The DPoE ONU MUST support classification based on one or more fields in an I-Tag.
- The DPoE ONU MUST support classification based on one or more fields in an S-Tag.
- The DPoE ONU MUST support classification based on one or more fields in a C-Tag.
- The DPoE ONU MUST support classification based on one or more fields in an S-Tag and one or more fields in a C-Tag simultaneously.
- The DPoE ONU SHOULD support classification based on other combinations of classifiers operating on fields of individual tags listed above.

When a TPID classifier is provisioned in a configuration file, it is expected that the corresponding VID classifier is also present Annex C.

When L2 classifiers are provisioned in the CM configuration file for an individual SF, the DPoE System and DPoE ONU adhere to the following rules:

- If one or more fields from a single tag (e.g., C-TPID or C-VID within a C-Tag) are provisioned in a classifier, the DPoE System MUST configure the D-ONU to match these fields against the outermost tag on a frame.
- If one or more fields from both an S-Tag and C-Tag (e.g., S-VID within S-Tag and C-VID within C-Tag) are provisioned in a classifier, the DPoE System MUST configure the D-ONU to match these parameters with the S-Tag as the outer tag and C-Tag as the inner tag.

¹⁶ Revised per MULPIv1.0-N-12.0022-1 and per MULPIv1.0-N-12.0030-1 on 5/29/12 by JB. Revised per MULPIv1.0-N-12.0042-3 on 8/16/12 by JB. Revised per MULPIv1.0-N-14.0129-2 on 3/6/14 by JB.

- If one or more fields from an I-Tag (e.g., I-SID within I-Tag) are provisioned in a classifier, the DPoE System MUST configure the D-ONU to match these parameters with or without the existence of the B-Tag in the frame.

7.3.4.1 Types of Classifiers¹⁷

The DPoE System supports Classifiers for MEF and IP(HSD) services. The DPoE ONU supports Classifiers for MEF and IP(HSD) services.

The set of classifiers used by the DPoE Network, applicable to both IP(HSD) and MEF SFs, is as defined in [MULPIv3.0] with the following exceptions:

- The "IP Classification Parameters" apply to IPv4 fields only. In the current version of this specification, IP Classification Parameters do not apply to IPv6.
- The vCM MUST use information provided in CM Interface Mask (CMIM) TLVs to map classifiers (and by reference, service flows) to individual DPoE ONU interfaces. The vCM MUST populate the docsQoS_pktClassCmInterfaceMask object in the docsQoS_pktClassTable for each DPoE ONU interface. In this way, the CMIM information is used for classification purposes, since the vCM translates CMIM into Object Context TLVs (0xD6 branch, see [DPoE-SP-OAMv1.0] for more details).
- Classifiers cannot be added to the frame classification table using the DOCSIS MAC sublayer service interface, which is not instantiated within the DPoE Network.
 - DPoE Classifiers do not support deferred activation, i.e., once the Classifier is provisioned, it becomes automatically activated.

The DPoE System SHOULD NOT use the EPON MAC address as a QoS classifier. The DPoE ONU SHOULD NOT use the EPON MAC address as a classifier.

[DPoE-SP-MEFv1.0] requires support for the following QoS classifiers fields:

- A set of [802.1ad] defined fields: C VLAN ID (C-VID), S VLAN ID (S-VID), C VLAN Type ID (C-TPID), S VLAN Type ID (S-TPID).
- Service Identifier (I-SID), I-SID Type ID (I-TPID), Backbone MAC Destination Address (B-DA). The DPoE System MUST support classification based on the Backbone VLAN ID (B-TCI) and B-SID Type ID (B-TPID).¹⁸

7.3.5 Service Classes¹⁹

The DPoE System MUST support configuration of QoS Parameters via Named Service Classes as specified in [MULPIv3.0], section 7.5.3. This functionality allows for configuration of a common set of QoS Parameters for SFs through the access to the "CMTS MIBs" as opposed to configuring them via CM configuration files. However, the CM configuration file still needs to refer to a specific Service Class Name configured on DPoE System.

7.3.6 Authorization

In the DPoE Network, the SF authorization process is not required, because the admission control and resource availability decisions are performed on the DPoE System, and the DPoE ONU does not participate in this process. The DPoE ONU is controlled by the DPoE System using the master-slave relationship.

7.3.7 SF and Classifiers²⁰

Classifiers associate packets into exactly one SF. The SF encodings provide the QoS Parameters for treatment of those packets on the EPON. Service Flow encodings are described in Section 7.3.1 and Annex C.

¹⁷ Revised per MULPIv1.0-N-12.0058-1 on 2/7/13 by JB.

¹⁸ This version of the specifications does not provide a TLV for a configuration file to provide B-Tag (B-TPID and B-TCI) settings to a DPoE Network. Without TLVs, this requirement can be met with CLI, XML, or other means of configuration that are not specified.

¹⁹ Revised per MULPIv1.0-N-12.0022-1 on 5/29/12 and MULPIv1.0-N-14.0117-2 on 2/18/14 by JB.

²⁰ Revised per MULPIv1.0-N-12.0010-1 on 3/27/12 by JB.

The definition of a primary SF in DPoE specifications is different from that in [MULPIv3.0]. A DPoE System supports CM configuration files, with and without primary SFs. If a primary service flow is identified it applies to the first UNI (lowest CMIM interface). In DPoE, a primary SF is not identified by its position in the CM configuration file, but rather based on the presence (or lack) of associated classifier entry in the configuration TLV. In this way, if all SFs present in the CM configuration file have an associated classifier, then the DPoE System does not support primary SF and the DPoE System operates without primary SFs for that DPoE ONU. If there are one or more SFs present in the CM configuration file that do not have an associated classification entry in the configuration TLVs, the DPoE System MUST treat the first such SF as primary, independent of its position in the CM configuration file, all the other SFs without an associated classifier are ignored.

Support for primary SFs in the DPoE specifications is introduced to maintain partial backward compatibility with the DOCSIS provisioning systems of some operators, who rely on the presence and support of primary SF for the implementation of IP(HSD) applications.

If all of the SFs in the CM configuration file are defined with associated classifiers, then the DPoE System MUST operate without primary SFs. The DPoE System MUST NOT require the existence of a primary downstream SF to perform downstream forwarding. The DPoE ONU MUST NOT require the existence of a primary upstream SF to perform upstream forwarding. When there are no primary SFs configured the DPoE System MUST drop all unclassified traffic. The DPoE System MUST set the docsQoSServiceFlowPrimary value to false for all SFs when operating without primary SFs.

If the CM configuration file contains at least one upstream SF without associated classifiers, the DPoE System MUST use the first such SF without an associated classifier entry in the configuration TLV as the primary upstream SF. If configured by the DPoE System, the DPoE ONU MUST use the primary upstream SF as the default upstream SF to forward all unclassified traffic.

If the CM configuration file contains at least one downstream SF without associated classifiers, the DPoE System MUST use the first such SF without an associated classifier entry in the configuration TLV as the primary downstream SF. The DPoE System MUST use the primary downstream SF as the default downstream SF for all unclassified downstream traffic. The DPoE System MUST set the docsQoSServiceFlowPrimary value to true for all primary SFs.

Classification of "DOCSIS MAC Management Messages", via the "Ethertype/DSAP/MacType" TLV, does not apply to DPoE Networks. "MAC Management Messages" in a DPoE Network are [802.3] OAM and OAM Extensions [DPoE-SP-OAMv1.0] and are not classifiable.

Policy-based classification via [RFC 2669] or other methods are not required for DPoE Networks.

7.4 Data Link Encryption Support

The procedures to support data-link encryption are defined in [DPoE-SP-SECv1.0]. The interaction between the MAC layer and the security system is limited to the items defined below.

7.4.1 MAC Messages

Encryption of the OAMPDUs is defined in [DPoE-SP-SECv1.0].

Encryption of the MPCPDUs is defined in [DPoE-SP-SECv1.0].

Encryption of the data bearing MAC DATAPDUs is defined in [DPoE-SP-SECv1.0].

7.4.2 Framing

Detailed information on the MAC framing format, when encrypted, is included in the [DPoE-SP-SECv1.0] specification.

8 DATA FORWARDING²¹

The basic architecture for DPoE Networks is described in [DPoE-SP-ARCHv1.0]. For MEF Services, the DPoE Network uses EPON technology to provide P2P Ethernet connections between the DPoE System and a DPoE ONU. P2P Ethernet connections in the form of Ethernet Virtual Connections (EVCs) are the underlying transport technology for Metro Ethernet services in the DPoE Network. IP(HSD) is a DPoE service that transports IP packets in the same manner as traditional DOCSIS and does not require the use of EVCs and [802.1ad] tagging across the TU interface.

This version of the DPoE specifications only supports P2P services and does not provide a standard method for implementing P2MP services, such as Ethernet broadcast, IP broadcast, IP multicast, IP anycast, etc. Although these services can be implemented "over the top" of the P2P services, such an implementation would (a) not take advantage of the P2MP (broadcast) capabilities of EPON, and (b) would therefore require duplication of frames (or packets for IP services) prior to forwarding on the DPoE Network. For example, Ethernet broadcast could be implemented using an external Ethernet bridge, but the DPoE System does not provide this capability in this version of DPoE specifications.

Since the DPoE System relies on a common set of Ethernet forwarding principles, all traffic (downstream or upstream) not classified (for forwarding) MUST be dropped at or before the applicable classification reference point.

8.1 MEF Forwarding Requirements

Under [802.1ad], the term S-VLAN typically refers to a Service (S) VLAN, and the term C-VLAN typically refers to a Customer (C) VLAN. In telecommunications access services implemented with DPoE Networks, for example, the S-VLAN and C-VLAN are used as "outer" and "inner" tags (q-in-q) without respect to their typical meaning in [802.1ad]. See [DPoE-SP-MEFv1.0] which provides a detailed explanation of the Metro Ethernet services.

There are two Ethernet forwarding models in DPoE Networks for Metro Ethernet services. These are the Provider Bridge (PB) and Provider Backbone Bridge (PBB). In the DPoE Network, the classifier is used for the purpose of forwarding frames. This includes making decisions on which frames are dropped or forwarded to specific LLIDs, depending on the set of provisioned classification rules.

8.1.1 Provider Bridge (PB)

The common model is a Metro Ethernet Forum service model that MUST be implemented with [802.1ad] by the DPoE System. The Metro Ethernet Forum service model MUST be implemented with [802.1ad] in the DPoE ONU. Data forwarding for each service is accomplished by establishing, at the time of provisioning, a globally unique combination of IEEE 802 standard Ethernet headers and the port information configured through the Cable Modem Interface Mask (CMIM).

PB forwarding by the DPoE System MUST use the following parameters. PB forwarding by the DPoE ONU MUST use the following parameters:

- S-TPID
- S-VLAN ID
- C-TPID
- C-VLAN ID
- C-DA
- C-SA

The DPoE Classifier (on the DPoE System (C_S) for downstream traffic and on the DPoE ONU (C_O) for upstream traffic) classifies (forwards) each EVC into an LLID.

²¹ Revised per MULPIv1.0-N-14.0117-2 on 2/19/14 by JB.

The DPoE System MUST default to a TPID of 0x88a8 as specified in [802.1ad], unless otherwise configured. The DPoE ONU MUST default to a TPID of 0x88a8 as specified in [802.1ad], unless otherwise configured. [DPoE-SP-MEFv1.0] offers optional methods of operator configurable TPID translation.

8.1.2 Provider Backbone Bridge (PBB)

The DPoE System MUST support [802.1ah] forwarding and classification. The DPoE ONU MUST support [802.1ah] forwarding and classification. [802.1ah] encapsulation is optional as detailed in [DPoE-SP-MEFv1.0].

Every DPoE System functions as an I-BEB (for downstream forwarding) as specified in [DPoE-SP-MEFv1.0]. In the upstream, DPoE System forwards PBB traffic based on provisioned [802.1ah] parameters, utilizing at least one of those identified below.

The DPoE System MAY use the following [802.1ah] frame fields to classify (forward) frames into individual SFs. The DPoE ONU MAY use the following [802.1ah] frame fields to classify (forward) frames into individual SFs.

- B-DA
- B-SA
- B-TPID²²
- B-TCI²³
- I-TPID
- I-SID

The DPoE System MUST default to a TPID of 0x88e7 as specified in [802.1ah], unless otherwise configured. The DPoE ONU MUST default to a TPID of 0x88e7 as specified in [802.1ah], unless otherwise configured. [DPoE-SP-MEFv1.0] offers optional methods of operator-configurable TPID translation.

8.2 Multicast Forwarding

IP multicast is not supported in this version of the DPoE specification. IP multicast services may be implemented "over the top" of DPoE IP(HSD) services or over DPoE Metro Ethernet services, but the current version of the DPoE specification does not specify the requirements for IP multicast for either IPv4 or IPv6. In either case, the support for the IP multicast services must not require any changes to this version of the DPoE specifications.

8.3 Requirements for IP(HSD) Forwarding

The DPoE System MUST support forwarding IPv4 frames that belong to an IP(HSD) service, using a combination of [802.1d] bridging and layer 3 routing. Within the DPoE System there are two modes of transmitting frames from the OLT to the Router. If the IP Serving Group associated with the upstream service flow specifies an S-VID the DPoE System OLT MUST add an S-Tag using the specified S-VID to the IP(HSD) frame. The DPoE System OLT MUST then transmit the frame to the Router. If the IP Serving Group associated with the upstream service flow does not specify an S-VID, the DPoE System OLT does not add an S-Tag and transmits the frame to the Router without further modification.

The DPoE ONU MUST support forwarding IP(HSD) frames having an IPv4 format. The DPoE ONU is intended to forward upstream IP(HSD) frames using 802.1d bridging. The DPoE ONU uses the combination of dynamic MAC address learning and static MAC provisioning to properly select the UNI port for downstream frames.

²² This version of the specifications does not provide a TLV for a configuration file to provide B-Tag (B-TPID and B-TCI) settings to a DPoE Network. Without TLVs, this requirement can be met with CLI, XML, or other means of configuration that is not specified.

²³ *ibid.*

8.3.1 IP Serving Group (IP-SG)

An IP Serving Group (IP-SG) represents a collection of one or more IP(HSD) upstream service flows, whose traffic is associated with a DPoE System physical or virtual router interface. The configuration of an IP-SG is specified in [DPoE-SP-IPNEv1.0]. The IP(HSD) frames are associated with an IP-SG based on the service flow attribute mask parameter configured in both the associated upstream service flow and the IP-SG. If the upstream service flow is not associated with an IP-SG, the DPoE System MUST use the default IP-SG configured .

9 DPOE SYSTEM AND DPOE ONU INTERACTION

9.1 DPoE ONU and vCM Initialization and Reinitialization

9.1.1 Scan for Downstream Channel

A DPoE ONU, upon initialization, by default, is not allowed to transmit until it is discovered by the DPoE System, ranged, registered, and granted for access. When powered on, it starts receiving the downstream channel data stream to (1) acquire and align the receive path clock to the data clock retrieved from the continuous transmission, and (2) synchronize its data path to the incoming frames.

Detailed description of the DPoE ONU synchronization process in 10G-EPON is included in 76.3.3.2 in [802.3].

DPoE ONU synchronization process for 1G-EPON is based on the legacy 1000BASE-X synchronization procedure, as described in 36.2.5.2.6 in [802.3]. Relationship between the legacy 1000BASE-X Physical Coding Sublayer (PCS) functions and EPON stack is depicted in Figure 65–4 in [802.3].

Since multi-channel operation is not supported in the DPoE System, there is also no need to store any last operational parameters in a non-volatile storage. Every time a DPoE ONU is initialized, it is required to go through a complete initialization, discovery, ranging, registration, and granting process until it is fully functional. There may be exceptions to this rule in case of support of advanced power-saving modes, where the DPoE ONU may have selectively disabled receive and/or transmit paths. In this case, it may be required to go only through the process of synchronization to the downstream channel and clock re-acquisition, while completely skipping the initialization, discovery, ranging, and registration steps. Specific details are outside of the scope of the DPoE specifications and are considered vendor-specific.

9.1.2 Continue Downstream Scanning

A DPoE ONU MUST NOT select any other transmission channel apart from the channel meeting the requirements of the [DPoE-SP-PHYv1.0] specification. That means that any functions related to channel scanning are not supported in the DPoE System.

9.1.3 Service Group Discovery and Initial Ranging

The DPoE System MUST NOT determine the service group of a DOCSIS CM for channel bonding and load balancing, since such functions are not supported by the underlying EPON transport layer.

Ranging in the DPoE System is performed first during the Discovery phase of the DPoE ONU registration, and then is carried out continuously during the regular operation, guaranteeing that the DPoE ONU RTT always remains correct and up-to-date. The ranging process for 10G-EPON is defined in 77.2.1.1 in [802.3], and for 1G-EPON is defined in 64.2.1.1 in [802.3].

DOCSIS-specific service group discovery and initial ranging functions are not supported in the DPoE System.

See 77.3.3 in [802.3] for 10G-EPON and 64.3.3 in [802.3] for 1G-EPON, for detailed information on the Discovery phase and related processes taking place on the DPoE System and DPoE ONU side. DPoE ONU Discovery represents the time when the DPoE ONU has successfully completed DPoE Network registration (see 64.3.3 in [802.3] and 77.3.3 in [802.3]) for some number of LLIDs (EPON logical links).

9.1.4 Authentication

See [DPoE-SP-SECv1.0] for details.

9.1.5 Establishing IP Connectivity

The DPoE ONU does not contain an IP stack and is not directly addressable using IP. The DPoE System MUST obtain an IP address on behalf of the DPoE ONU. The DPoE System performs IP provisioning for each DPoE ONU as they are discovered across the TU interface. The DPoE System MUST maintain the IP address and associated parameters for the DPoE ONU as a vCM.

Upon successful completion of IP Address assignment, the DPoE System obtains the DOCSIS CM configuration file via TFTP, followed by the processing of the DOCSIS configuration file. The diagram in Figure 6 shows an overview of the DPoE System establishing IP Connectivity, which takes place for each discovered DPoE ONU.

The DPoE provisioning is almost identical to what is specified in previous versions of [MULPIv3.0]. This section specifies only requirements that are new or different from the previous DOCSIS MULPI specifications with respect to establishment of IP Connectivity within a DPoE System. The DPoE System performing IP provisioning on behalf of a DPoE ONU MUST follow the operational flow of Figure 6.

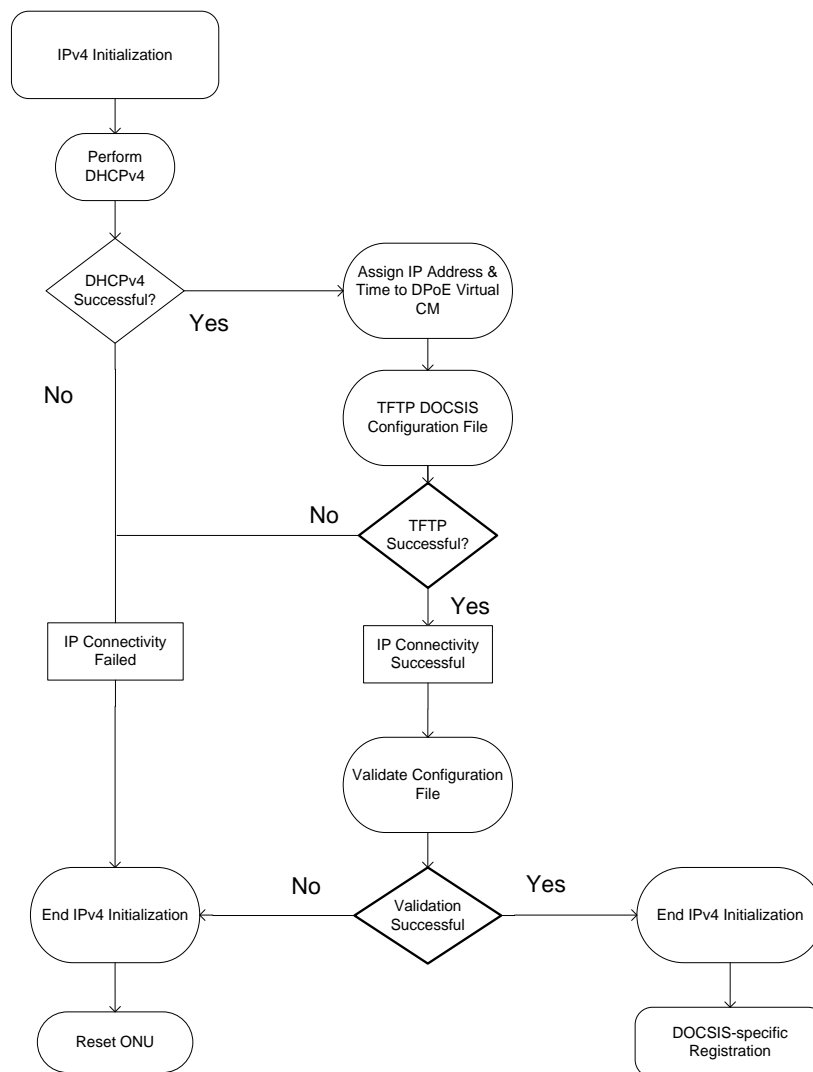


Figure 6 - IPv4 Only Provisioning Mode for DOCSIS-specific Registration

9.1.5.1 Establish IPv4 Network Connectivity²⁴

This section describes how the DPoE System provisions an IPv4 address and associated parameters on behalf of a DPoE ONU. The DPoE System MUST support the requirements as specified in [MULPIv3.0] specification. Figure 7 shows the DHCPv4 message sequence for a DPoE System. Establishment of IPv4 Network Connectivity by the DPoE System MUST take place after the discovery of the DPoE ONU.

²⁴ Revised per MULPIv1.0-N-12.0030-1 on 5/29/12 by JB.

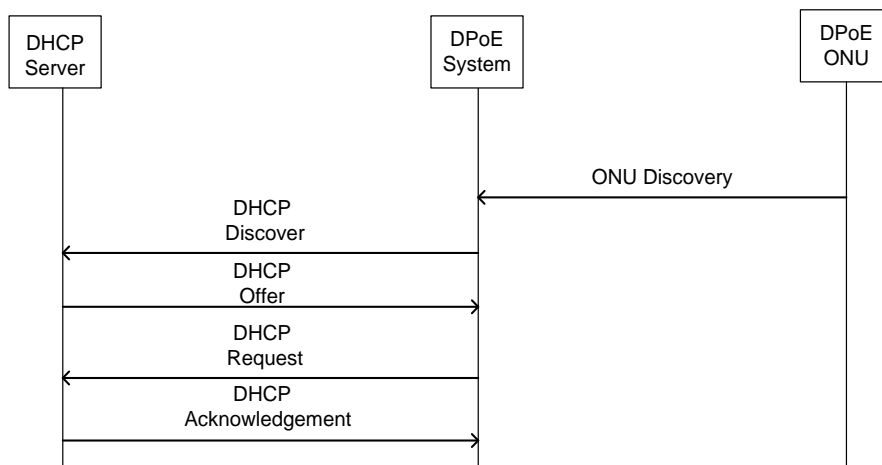


Figure 7 - Establishment of IPv4 Network Connectivity

As specified in [DPoE-SP-OSSIV1.0], the DPoE System MUST maintain the appropriate DOCSIS state for each vCM known to the DPoE System. This is the DPoE System view of the vCM (docsIf3CmtsCmRegStatusValue of the DOCS-IF3-MIB).

- When the DPoE System generates a DHCPv4 Discover message to obtain an IP Address for the vCM, the DPoE System MUST transition the vCM to the "startDhcpv4" state.
- Upon successful completion of IPv4 Address assignment for the vCM, the DPoE System MUST transition the vCM to the "dhcpv4Complete" state.
- If the DHCPv4 Address assignment fails, the DPoE System MUST reset the DPoE ONU and transition the vCM to the "other" state.

As specified in [DPoE-SP-OSSIV1.0], the DPoE System MUST maintain the appropriate DOCSIS state on behalf of the vCM. This is the vCM view of its own internal state (docsIf3CmStatusValue of the DOCS-IF3-MIB).

- When the DPoE System generates a DHCPv4 Discover message to obtain an IP Address for the vCM, the DPoE System MUST transition the vCM to the "dhcpv4InProgress" state.
- Upon successful completion of IPv4 Address assignment for the vCM, the DPoE System MUST transition the vCM to the "dhcpv4Complete" state.
- If DHCPv4 Address assignment fails, the DPoE System MUST reset the DPoE ONU and transition the vCM to the "other" state.

A vCM with an unexpired IPv4 address MUST send a DHCPRELEASE message as described in [RFC 2131] immediately prior to a reset caused by a set to the docsDevResetNow attribute.

9.1.5.1.1 DHCPv4 Fields Used by the vCM²⁵

The DPoE System MUST support the DHCP requirements for this section as specified in [MULPIv3.0], with the following exceptions.

The following fields MUST be present in the DHCPDISCOVER and DHCPREQUEST messages from the DPoE System:

- The client hardware address (chaddr) is set to the 48-bit EPON MAC address of the DPoE ONU.
- The parameter request list option is included. The option codes are defined in [RFC 2132] and [RFC 4361]. The following option codes MAY be included in the list: Option code 1 (Subnet Mask), Option code 2 (Time

²⁵ Revised per MULPIv1.0-N-12.0017-2 and MULPIv1.0-N-12-0030-1 on 5/29/12 by JB.

Offset), Option code 3 (Router Option), Option code 4 (Time Server Option) The following option codes MUST be included in the list: Option code 7 (Log Server Option) and Option code 125 (DHCPv4 Vendor-Identifying Vendor Specific Information Option). Option code 125 (DHCPv4 Vendor-Identifying Vendor Specific Information Options) for DOCSIS 3.0 defined in [CANN-DHCP-Reg], with the following sub-options: Sub-option code 1, the DHCPv4 Option Request option. The following option code MUST be included in the DHCPv4 Option Request option: Sub-option code 2, DHCPv4 TFTP Servers Option.

- Option code 60 (Vendor Class Identifier) — the following ASCII-encoded string is the value to be present in Option code 60: docsis3.0.

The following fields MAY be present in the DHCPDISCOVER and DHCPREQUEST messages from the DPoE System:

- Option code 125 (DHCPv4 Vendor- Identifying Vendor Specific Information Option): Sub-option code 5, Modem Capabilities Encoding for DHCPv4

Sub-Option code 5 SHOULD NOT be supported by the DPoE System in this version of the DPoE specification.

9.1.5.1.2 Use of T1 and T2 Timers

Since the DPoE System performs DHCP on behalf of a vCM, it is the DPoE System that MUST comply with the DHCP T1/T2 requirements defined in [MULPIv3.0].

9.1.5.1.3 DHCPv4 Renew Fields Used by the vCM

It is possible, during the DHCPv4 renew operation, that the vCM will receive updated fields in the DHCPACK message.

If any of the IP address (yiaddr), the Subnet Mask, or the Next Hop Router (router option) are different in the DHCPACK than the current values used by the DPoE System for the vCM, the DPoE System MUST do one of the following:

- Reinitialize the associated DPoE ONU.
- Change the vCM's addressing to use the new values without reinitializing the associated DPoE ONU.

If the Config File Name or the SYSLOG server address is different in the DHCPACK than the current values used by the DPoE System for the vCM, the DPoE System MUST ignore the new fields.

If the Time Offset or Time server address values are different in the DHCPACK than the current values used by the DPoE System for the vCM, the DPoE System will take no action, as these fields do not apply for the DPoE System.

9.1.5.1.4 DPoE System Requirements²⁶

The DPoE System MUST support DHCPv4 Option 43 sub-options identified in [CANN-DHCP-Reg] and [eDOCSIS]. The DPoE System MUST support the DHCPv4 Relay Agent requirements as specified in [MULPIv3.0], with the following exceptions.

In order to assist the DHCP server in differentiating between a DHCPDISCOVER sent from a vCM and a DHCPDISCOVER sent from a CPE:

- The DPoE System DHCPv4 Relay Agent MUST support the DHCP Relay Agent Information Option (RAIO), according to [RFC 3046] Specifically, the DPoE System DHCPv4 Relay Agent MUST add a RAIO to the DHCPDISCOVER message before relaying the message to a DHCP server. The DPoE System MUST include in the RAIO, the 48-bit MAC address of the DPoE ONU for the DHCPDISCOVER in the agent remote ID sub-option field [RFC 3046] The DPoE System MUST support "DHCPv4 Relay Agent CMTS capabilities option" as specified in [CANN-DHCP-Reg] including the following sub-options: 'CMTS DOCSIS Version Number', 'DPoE System Version Number', 'DPoE System DHCPv4 PBB service option', 'CMTS CM Service Class', and 'CMTS MSO Defined Text'.
 - The DPoE System MUST set the 'CMTS DOCSIS Version Number' sub-option to "3.0".
 - The DPoE System MUST set the 'DPoE System Version Number' sub-option to "1.0".

²⁶ Revised per MULPIv1.0-N-12.0017-2 on 5/29/12 by JB.

9.1.5.2 Establish Time of Day

The DPoE System MUST establish a time reference through Time of Day (TOD), Network Time Protocol (NTP), or other means prior to establishing vCMs. The DPoE System is not required to establish this time reference for each vCM, but the DPoE System MUST supply a time reference for each vCM based on the already established DPoE System time. This time reference MUST be used by each vCM in the DPoE System for a number of purposes, including, among the others, time stamping in the collected log files. The use cases for time references are vendor-specific and outside the scope of this specification.

9.1.5.3 Transfer of Operational Parameters

The DPoE System MUST obtain the CM configuration file for the applicable DPoE ONU as specified in [MULPIv3.0] for IPv4 with no exceptions.

As specified in [DPoE-SP-OSSIV1.0], the DPoE System MUST maintain appropriate DOCSIS state for each vCM known to the DPoE System. This is the CMTS view of the CM (docsIf3CmtsCmRegStatusValue of the DOCS-IF3-MIB).

- When the DPoE System generates a TFTP request to obtain the DOCSIS CM configuration file, the DPoE System MUST transition the vCM to the "startConfigFileDownload" state.
- Upon successful TFTP of the DOCSIS CM configuration file for the vCM, the DPoE System MUST transition the vCM to the "configFileDownloadComplete" state.
- If the TFTP of the DOCSIS CM configuration file fails, the DPoE System MUST reset the DPoE ONU and transition the vCM to the "other" state.

As specified in [DPoE-SP-OSSIV1.0], the DPoE System MUST maintain appropriate DOCSIS state on behalf of the vCM. This is the CM view of its own, internal state (docsIf3CmStatusValue of the DOCS-IF3-MIB).

- Upon successful TFTP of the DOCSIS CM configuration file for the vCM, the DPoE System MUST transition the vCM to the "configFileDownloadComplete" state.
- If TFTP of the DOCSIS CM configuration file fails, the DPoE System MUST reset the DPoE ONU and transition the vCM to the "other" state.

9.1.5.4 Configuration File Processing²⁷

After downloading the configuration file, and prior to commencing the OAM Provisioning process, the DPoE System performs several processing steps with the configuration file on behalf of DPoE ONU.

The DPoE System MUST support the TLV encodings specified in Annex C of this document. The DPoE System MUST perform TLV validation according to Annex C in [MULPIv3.0]. Examples of such validation include TLV range checking, parameter applicability, and parameter interdependency checks.

The DPoE System performs operations to verify the validity of a configuration file and MUST reject a configuration file that is invalid. An invalid configuration file has any of the following characteristics:

- Lacks one or more mandatory items, as defined in this specification and [MULPIv3.0].
- Has one or more TLV-11 encodings that cannot be processed and cause rejection of the file.
- Contains a TLV-54 encoding, SNMPv3 Access View Configuration, that causes rejection of the file, as defined in [MULPIv3.0].
- Contains specifications for a number of SFs that cannot be supported due to the number of EPON Links registered by the DPoE ONU.

The DPoE System MUST reject a configuration file if it is considered as invalid under the conditions specified above. The DPoE System MAY reject a configuration file that has an invalid CM MIC, CMTS MIC or Extended CMTS MIC as defined in [MULPIv3.0]. The DPoE System MAY continue with registration under conditions other than specified above.

²⁷ Revised per MULPIv1.0-N-12.0011-1 on 3/27/12 by JB.

If the DOCSIS CM configuration file validation fails, the DPoE System MUST reset the DPoE ONU and transition the vCM view (docsIf3CmStatusValue) to the "accessDenied" state, and the DPoE System view of the vCM (docsIf3CmtsCmRegStatusValue) to the "other" state (as defined in [DPoE-SP-OSSIv1.0]).

9.1.5.5 Post-Registration Failures to Renew IP Address

The DPoE System and vCM MUST support the Post-registration Failures and renew IPv4 Address requirements as specified in [MULPIv3.0] with no exceptions.

9.1.6 Registration with the DPoE System

Once the CM Configuration file is validated, the DPoE System performs registration. The DOCSIS-specified Registration process of Registration Request, Response, and Acknowledgement between the DPoE System and vCM is internal processing by the DPoE System, and thus is not mandated or specified by this document.

Figure 8 shows the DPoE System Registration process.

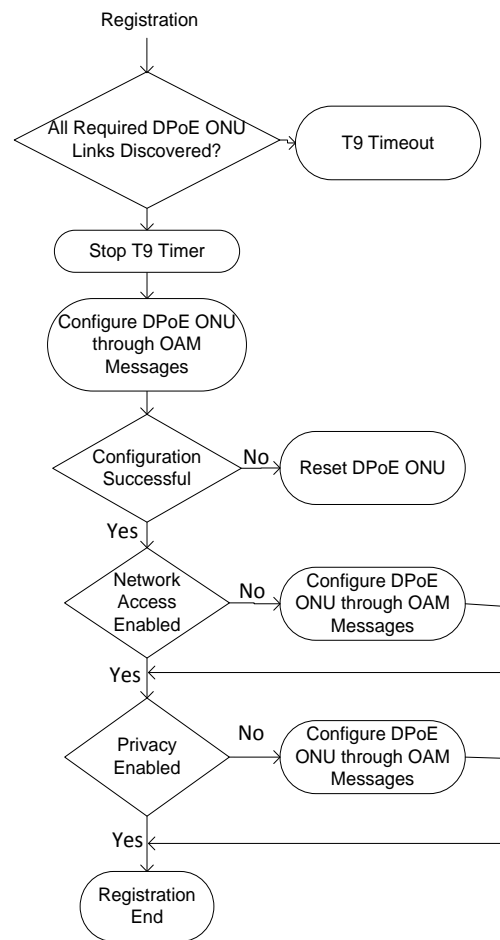


Figure 8 - Registration of DPoE ONU in DPoE System

9.1.6.1 DPoE System vCM Requirements

The DPoE System MUST transition the state of the vCM (docsIf3CmStatusValue) to "operational" state upon completion of successful registration. Successful registration is defined as successful completion of processing of the CM configuration file by the DPoE System and the successful configuration of the DPoE ONU using the DPoE OAM messages.

The DPoE System MUST NOT transition the state of the vCM (docsIf3CmStatusValue) to the "operational" state upon completion of unsuccessful registration. Unsuccessful registration is defined as the failed processing of the CM configuration file by the DPoE System, or the unsuccessful configuration of the DPoE ONU using the DPoE OAM messages. Such failures MUST result in the vCM state in the DPoE System being set to "accessDenied" (docsIf3CmStatusValue).

If the DOCSIS CM configuration file specifies that network access is disabled, the DPoE System MUST configure the DPoE ONU with the parameters specified in the DOCSIS CM configuration file, and block all frame forwarding from S interfaces on the DPoE ONU. The DPoE System MUST also transition the vCM state (docsIf3CmStatusValue) to "forwardingDisabled".

The DPoE System MUST NOT allow the transmission of customer data until the successful registration is complete.

9.1.6.2 DPoE System Requirements

Upon the successful validation of the DOCSIS CM configuration file, the DPoE System MUST transition the vCM state (docsIf3CmtsCmRegStatusValue) to "startRegistration". During registration, the DPoE System MUST configure both itself and the DPoE ONU based on the TLVs specified in the DOCSIS CM configuration file. The DPoE System MUST configure the DPoE ONU using the appropriate OAM messages.

Any failures during the configuration of the DPoE System or DPoE ONU MUST result in the DPoE System resetting the DPoE ONU and transitioning the DPoE System view of the vCM (docsIf3CmtsCmRegStatusValue) to the "other" state.

During the configuration of the DPoE ONU, all of the required DPoE ONU Links MUST be discovered during the T9 timeout period by the DPoE System. If the T9 timer expires, the DPoE System MUST reset the DPoE ONU and transition the vCM (docsIf3CmtsCmRegStatusValue) to the "other" state.

If the DOCSIS CM configuration file specifies that network access is disabled, the DPoE System MUST configure both the DPoE System and DPoE ONU with the parameters specified in the DOCSIS CM configuration file, and block all frame forwarding from S interfaces on the DPoE ONU. The DPoE System MUST also transition the vCM state (docsIf3CmtsCmRegStatusValue) to "forwardingDisabled".

If the DOCSIS CM configuration file specifies that privacy is disabled, the DPoE System MUST disable encryption on all of the LLIDs (EPON logical links) of the DPoE ONU being used to carry DOCSIS SFs.

Upon successful completion of registration processing, the DPoE System MUST transition the vCM state (docsIf3CmtsCmRegStatusValue) to "operational" for the DPoE ONU. The DPoE System MUST NOT allow the transmission of data until successful registration is complete.

9.1.7 Service IDs During vCM Initialization

A primary SID MUST be reserved and assigned to a vCM when it is instantiated by the DPoE System. Additional SIDs (for secondary upstream flows) MUST be reserved and assigned after TFTP, prior to OAM provisioning by the DPoE System.

9.2 Periodic Maintenance

In the DPoE Network, periodic maintenance of the underlying EPON transport layer is limited to continuous ranging (see 77.2.1.1 in [802.3av] for 10G-EPON and 64.2.1.1 in [802.3] for 1G-EPON definitions), based on the MPCP GATE and REPORT MPCPDUs, exchanged between the DPoE System and connected DPoE ONUs, and carrying a timestamp field. This particular field is used to calculate the RTT for the given DPoE ONU in a continuous manner, allowing the DPoE System to dynamically adjust to any changes in the path delay without the need to bring the DPoE ONU down for re-registration. In this way, the EPON system guarantees timely and accurate information on the DPoE ONU specific path delay without affecting the carried services.

Additionally, each DPoE System has the ability to monitor continuously the status of the LLIDs associated with the given DPoE ONU, keeping track of the responses generated by each LLID. A keep-alive mechanism is implemented in EPON, where the DPoE System maintains a watchdog for each registered LLID. Upon expiration of the said watchdog, the associated LLID is considered inactive (for whatever reason) and the given logical link is

deregistered. A DPoE ONU with at least one LLID that was flagged as inactive **MUST** be considered as inactive by the DPoE System, affecting in this way all the LLIDs associated with the given DPoE ONU.

In the EPON specific keep-alive mechanism, each LLID will periodically announce its activity to the DPoE System by sending a REPORT MPCPDU in the granted bandwidth slot. The REPORT MPCPDU may be solicited by the DPoE System (hence every time DPoE ONU is requested to send the REPORT MPCPDU with the current snapshot of its queues, this is considered by the DPoE System as confirmation that the given LLID is active) or unsolicited (in which case the DPoE ONU will automatically include the REPORT MPCPDU in the assigned transmission slot). Effectively, the DPoE System is notified on the LLID status and may track LLID activity for all registered DPoE ONUs in the system. Decision on whether to deregister a complete DPoE ONU (or not) if one of the logical links associated with this ONU is found to be inactive, is implementation-specific and out of the scope of the DPoE specifications.

Keep-alive mechanism is referred to in 77.3.6.2 in [802.3] for 10G-EPON and in 64.3.6.2 in [802.3] for 1G-EPON, but was never formally specified in the EPON standards.

Other DOCSIS-specific maintenance functions, including periodic ranging, are not supported in the DPoE System.

9.3 Fault Detection and Recovery

9.3.1 MAC Layer Error-Handling

This section describes the procedures that are required when an error occurs at the MAC framing level.

The most obvious type of error occurs when the FCS in the MAC frame fails. In such a case, Ethernet-specific handling is performed, per [802.3] definitions. Specifically, a corrupted Ethernet frame (where calculated FCS does not match the received FCS) is always dropped at the MAC layer to prevent error propagation across Ethernet links.

In the upstream channel, there are two possible types of errors, namely:

- Loss of burst synchronization, where the DPoE System receiver fails to synchronize to the incoming data burst synchronization fields, losing the ability to properly delineate individual frames carried in the burst. This type of error is less common in 1 Gbps links, where data rate is lower and the framing structure does not feature a mandatory stream-based FEC. In 10 Gbps links, inability to synchronize to the incoming burst headers, and specifically inability to acquire the Burst Delimiter (see 76.3.2.5.1 in [802.3], applicable to 10G-EPON only), will cause the DPoE System to lose a complete data burst. Luckily, 10G-EPON system is designed to operate correctly at BER as low as 10^{-3} , which guarantees proper system operation under very unfavorable transmission conditions.
- Burst overlap, which occurs when at least two DPoE ONUs transmit in such a way that their data bursts overlap (either partially or completely) at the DPoE System receiver. In such a case, the DPoE System may be able to recover part of the initial burst only. The rest of the overlapped data bursts are lost. EPON does not have a mechanism to signal burst collision to individual DPoE ONUs and request their retransmission. EPON relies for retransmission on higher layer protocols (e.g., TCP/IP). Under normal operating conditions, burst overlap is not expected to occur. Burst overlap is expected to occur only during the Discovery Window (see 77.3.3 in [802.3] for 10G-EPON and 64.3.3 in [802.3] for 1G-EPON).

Additionally, EPON features FEC functionality, as described in [DPoE-SP-PHYv1.0].

9.4 UNI Management²⁸

The configuration of the given UNI port on the D-ONU may be configured via the CM configuration file downloaded by the vCM from the TFTP server during the vCM initialization process, following the process outlined in section 9.1. The configuration of the UNI parameters such as auto-negotiation function, operating speed, duplex, admin status, Energy Efficient Ethernet (EEE), Power over Ethernet (PoE), as well as media type may be set to specific values at the time when the D-ONU is configured by the vCM at the vCM initialization time. The UNI parameters are defined using TLV 79 and its sub-TLVs as described in [MULPIv3.0].

²⁸ Revised per MULPIv1.0-N-14.0125-2 on 2/27/14 by JB.

The operator may further modify individual UNI parameters via CLI or SNMP, if needed, though any changes to these parameters are non-persistent; i.e., when the vCM is reset / reboots, configuration for individual UNI parameters is read from the downloaded CM configuration file.

10 DOWNLOADING CABLE MODEM OPERATING SOFTWARE

The protocol requirements are included in [DPoE-SP-OAMv1.0], with further requirements included in [MULPIv3.0] and [DPoE-SP-SECv1.0].

The internal structure of the software package, delineation process, etc., are vendor-specific and outside the scope of the DPoE specifications.

Annex A Well-Known Addresses

In the DPoE specifications, the list of well-known addresses is composed of definitions included in section A.1.1 of [MULPIv3.0], with the additional extensions as mentioned below.

- Multicast MAC Control address in EPON, used as MAC DA for MPCPDUs: 0x01-80-C2-00-00-01
- Slow Protocols Multicast address, used as MAC DA in DPoE OAMPDUs: 0x01-80-C2-00-00-02
- Specific ranges of LLIDs are also well-known, namely:
 - Broadcast LLID in 1G-EPON: 0x7F-FF
 - Broadcast LLID in 10G-EPON: 0x7F-FE
 - Reserved range of LLIDs in 10G-EPON: 0x7F-FD – 0x7F-00

Annex B Parameters and Constants²⁹

All EPON-specific parameters and constants are defined in the respective EPON standards, i.e., [802.3], Clause 64 and Clause 65 for 1G-EPON, and [802.3], Clause 77 for 10G-EPON. Their reproduction in this document is considered unnecessary.

The DPoE System SHOULD support the following parameters from Table B-1 of DOCSIS [MULPIv3.0] Annex B, as appropriate:

- TFTP Backoff Start
- TFTP Backoff End
- TFTP Request Retries
- TFTP Download Retries
- TFTP Wait
- T9 Timer

²⁹ Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

Annex C Common TLV Encodings

The DPoE System MUST parse and apply the provisioning parameters defined by the TLVs contained within the CM provisioning files as part of the CM registration process. Note that other TLVs are defined in Annex C of [MULPIv3.0], but they are not present in the provisioning file and, therefore, are not covered in this table.

The following tables, C.1 through C.12 enumerate whether support is required for a particular TLV by the DPoE System in this version of the DPoE specifications. Four columns are provided in the table:

- TLV – TLV Number as defined in Annex C of [MULPIv3.0]
- Name – Descriptive Name associated with the TLV
- Support Needed – indicating whether the support for the given TLV is needed or not.
- Comments – Additional information regarding limitations for the support of the TLV or an explanation as to why the TLV need not be supported.

When the DPoE System encounters a TLV that is not supported, then the DPoE System MUST ignore the TLV and allow the DPoE ONU to register normally, following the DPoE Network registration process.

C.1 [802.1ad] S-VLAN and C-VLAN Frame Classification Encodings³⁰

This field defines the parameters associated with [802.1ad] S-VLAN and C-VLAN frame classification.

Type	Length	Value
[22/23].14	n	

C.1.1 [802.1ad] S-VLAN TPID³¹

The values of the field specify the matching parameters for the [802.1ad] S-VLAN TPID field.

If this parameter is not specified for an entry, then the DPoE System MUST use a default value of 0x88a8 for the [802.1ad] S-VLAN TPID field. If this parameter is not specified for an entry, then the DPoE ONU MUST use a default value of 0x88a8 for the [802.1ad] S-VLAN TPID field. Other values of [802.1ad] S-VLAN TPID may be provisioned, as required.

The DPoE System MUST NOT match Ethernet frames without the [802.1ad] S-VLAN tag to this entry. The DPoE ONU MUST NOT match Ethernet frames without the [802.1ad] S-VLAN tag to this entry.

The S-VLAN TPID classifier is not intended to be used by itself. The DPoE System MUST reject a CM configuration file in which a S-VLAN TPID classifier is provisioned without a corresponding S-VLAN VID classifier.

Type	Length	Value
[22/23].14.1	2	stpid1, stpid2

No restriction on the range of stpid1, stpid2 is specified.

C.1.2 [802.1ad] S-VLAN VID³²

The values of the field specify the matching parameters for the [802.1ad] S-VLAN VID field. Only the least-significant 12 bits of the specified S-VLAN VID field are significant; the remaining four bits MUST be ignored for comparison by the DPoE System. The DPoE ONU MUST ignore the remaining four bits for comparison as only the least-significant 12 bits of the specified S-VLAN VID field are significant.

³⁰ Revised per MULPIv1.0-N-11.0005-1 on 5/16/11 by JB.

³¹ Revised per MULPIv1.0-N-12.0042-3 on 8/16/12 by JB.

³² Revised and per MULPIv1.0-N-12.0046-1 on 10/31/12 by JB.

If this parameter is specified for an entry, the DPoE System MUST NOT match Ethernet frames without [802.1ad] S-VLAN tag to this entry. If this parameter is specified for an entry, the DPoE ONU MUST NOT match Ethernet frames without the [802.1ad] S-VLAN tag to this entry.

Type	Length	Value
[22/23].14.2	2	svid1, svid2

No restriction on the range of svid1, svid2 is specified.

C.1.3 [802.1ad] C-VLAN TPID³³

The values of the field specify the matching parameters for the [802.1ad] C-VLAN TPID field.

If this parameter is not specified for an entry, then the DPoE System MUST use a default value of 0x8100 for the [802.1ad] C-VLAN TPID field. If this parameter is not specified for an entry, then the DPoE ONU MUST use a default value of 0x8100 for the [802.1ad] C-VLAN TPID field. Other values of [802.1ad] C-VLAN TPID may be provisioned, as required.

The DPoE System MUST NOT match Ethernet frames without the [802.1ad]C-VLAN tag to this entry. The DPoE ONU MUST NOT match Ethernet frames without the [802.1ad] C-VLAN tag to this entry.

The C-VLAN TPID classifier is not intended to be used by itself. The DPoE System MUST reject a CM configuration file in which a C-VLAN TPID classifier is provisioned without a corresponding C-VLAN VID classifier.

Type	Length	Value
[22/23].14.5	2	ctpid1, ctpid2

No restriction on the range of ctpid1, ctpid2 is specified.

C.1.4 [802.1ad] C-VLAN VID³⁴

The values of the field specify the matching parameters for the [802.1ad] C-VLAN VID field. Only the least-significant 12 bits of the specified C-VLAN VID field are significant; the remaining four bits MUST be ignored for comparison by the DPoE System. The DPoE ONU MUST ignore the remaining four bits for comparison as only the least-significant 12 bits of the specified C-VLAN VID field are significant.

If this parameter is specified for an entry, the DPoE System MUST NOT match Ethernet frames without [802.1ad] C-VLAN tag to this entry. If this parameter is specified for an entry, the DPoE ONU MUST NOT match Ethernet frames without the [802.1ad] C-VLAN tag to this entry.

The use of TLV [22/23].14.6 in the DPoE CM config files is preferred over the use of the legacy DOCSIS TLV [22/23].11.2. A DPoE CM config file should not use both TLV [22/23].11.2 and TLV [22/23].14.6 within a single instance of classifier. A DPoE System MUST use the value carried in the TLV [22/23].14.6, when a DPoE CM config file uses both TLV [22/23].11.2 and TLV [22/23].14.6 within a single instance of classifier.

Type	Length	Value
[22/23].14.6	2	cvid1, cvid2

No restriction on the range of cvid1, cvid2 is specified.

C.2 [802.1ah] I-TAG Packet Classification Encodings

This field defines the parameters associated with [802.1ah] I-TAG packet classification.

³³ Revised per MULPIv1.0-N-11.0005-1 on 5/16/11 and per MULPIv1.0-N-12.0042-3.on 8/16/12 by JB.

³⁴ Revised per MULPIv1.0-N-11.0005-1 on 5/16/11 and per MULPIv1.0-N-12.0046-1 on 10/31/12 by JB.

Type	Length	Value
[22/23].15	n	

C.2.1 [802.1ah] I-TAG I-TPID³⁵

The values of the field specify the matching parameters for the [802.1ah] I-TAG I-TPID field.

If this parameter is not specified for an entry, then the DPoE System MUST use a default value of 0x88e7 for the [802.1ah] I-TAG I-TPID field. If this parameter is not specified for an entry, then the DPoE ONU MUST use a default value of 0x88e7 for the [802.1ah] I-TAG I-TPID field. Other values of [802.1ah] I-TAG I-TPID may be provisioned, as required.

The DPoE System MUST NOT match Ethernet frames without the [802.1ah] I-TAG tag to this entry. The DPoE ONU MUST NOT match Ethernet frames without the [802.1ah] I-TAG tag to this entry.

The I-TAG I-TPID classifier is not intended to be used by itself. The DPoE System MUST reject a CM configuration file in which a I-TAG I-TPID classifier is provisioned without a corresponding I-TAG I-SID classifier.

Type	Length	Value
[22/23].15.1	2	itpid1, itpid2

No restriction on the range of itpid1, itpid2 is specified.

C.2.2 [802.1ah] I-TAG I-SID

The values of the field specify the matching parameters for the [802.1ah] I-TAG I-SID field.

If this parameter is specified for an entry, the DPoE System MUST NOT match Ethernet packets without the [802.1ah] I-TAG tag to this entry. If this parameter is specified for an entry, the DPoE ONU MUST NOT match Ethernet packets without the [802.1ah] I-TAG tag to this entry.

Type	Length	Value
[22/23].15.2	3	isid1, isid2, isid3

No restriction on the range of isid1, isid2, isid3 is specified.

C.3 Top Level TLVs³⁶

TLV	Name	Support Needed	Comments
0	Pad	MUST	
1	Downstream Frequency	MUST NOT	Not applicable, as there is only one downstream channel, the EPON itself.
2	Upstream Channel ID	MUST NOT	Not applicable, as there is only one upstream channel, the EPON itself.
3	Network Access Control Object	MUST	
4	DOCSIS 1.0 Class of Service	MUST NOT	No reason to carry this support forward to DPoE specification.
6	CM Message Integrity Check	MAY	The operator network is considered to be secure and this TLV is not needed as the DPoE System is getting the file directly from the TFTP Server. This TLV may be supported for backwards compatibility.

³⁵ Revised per MULPIv1.0-N- 2.0042-3.on 8/16/12 by JB.

³⁶ Revised per MULPIv1.0-N-12.0011-1 on 3/27/12 and per MULPIv1.0-N-12.0024-1 on 5/29/12, and revised again per MULPIv1.0-N-14.0138-1 on 3/6/14 by JB.

TLV	Name	Support Needed	Comments
7	CMTS Message Integrity Check	MAY	The operator network is considered to be secure and this TLV is not needed as the DPoE System is getting the file directly from the TFTP Server. This TLV may be supported for backwards compatibility
9	SW Upgrade Filename	MUST	
10	SNMP Write Access Control	SHOULD NOT	Operator feedback was that this TLV is not needed for this version of the specifications.
11	SNMP MIB Object	MUST	See the TLV 11 section for more details.
14	CPE Ethernet MAC Address	SHOULD NOT	Operator feedback was that this TLV is not needed for this version of the specifications.
17	Baseline Privacy	MUST NOT	Encryption is different on EPON, and configurable controls are limited. You can configure Traffic Key lifetime per-PON and per-Link. This doesn't correspond directly to TLV 17 items, however. See the Security (TLV 17) section for more details.
18	Max Number of CPEs	MUST	
19	TFTP Server Timestamp	MUST NOT	ONUs do not access the config file server, so there's no middle man or hackable entity like a CM to worry about.
20	TFTP Server Provisioned Modem IPv4 Address	MUST NOT	ONUs have IP addresses proxied by the DPoE System. Thus, this TLV has little value.
21	SW Upgrade IPv4 TFTP Server	MUST	
22	Upstream Packet Classification	MUST	See the Classification (TLVs 22 & 23) section for more details.
23	Downstream Packet Classification	MUST	See the Classification (TLVs 22 & 23) section for more details.
24	Upstream SF	MUST	See the Service Flows (TLVs 24 & 25) section for more details.
25	Downstream SF	MUST	See the Service Flows (TLVs 24 & 25) section for more details.
26	Payload Header Suppression	MUST NOT	EPON doesn't define header suppression.
28	Maximum Number of Classifiers	SHOULD NOT	This is admission control for classification resources.
29	Privacy Enable	MUST	This TLV enables/disables device certificate authentication and traffic encryption functions on the DPoE System and D-ONU. See [DPoE-SP-SECv1.0] for more detail. ³⁷
32	Manufacturer Code Verification Certificate	MUST	
33	Co-Signer Code Verification Certificate	MUST	
34	SNMPv3 Kickstart Value	SHOULD	
35	Subscriber Mgmt Control	MUST	
36	Subscriber Mgmt CPE IPv4 List	MUST	
37	Subscriber Mgmt Filter Groups	MUST	
38	SNMPv3 Notification Receiver	MUST	
39	Enable 2.0 Mode	MUST NOT	EPON PHY is not DOCSIS PHY.
40	Enable Test Modes	MUST NOT	Not applicable to EPON ONUs.
41	Downstream Channel List	MUST NOT	ONUs lock to a single DS PON channel.
42	Static Multicast MAC Address	SHOULD NOT	Multicast is not supported in this version of the specifications.

³⁷ Revised per MULPIv1.0-N-13.0083-1 on 6/19/13 by JB.

TLV	Name	Support Needed	Comments
43	DOCSIS Extension Field	MUST	See the TLV 43, [DPoE-SP-MEFv1.0] and [L2VPN] (TLVs 43.5, 45, & 65), and Customer (Subscriber) Management (TLV 43.7) sections for more detail.
45	Downstream Unencrypted Traffic (DUT) Filtering	SHOULD NOT	Not needed for this version of the specifications.
53	SNMPv1v2c Coexistence	MUST	
54	SNMPv3 Access View	MUST	
55	SNMP CPE Access Control	SHOULD NOT	There are currently no requirements to support this in this version of the specifications.
56	Channel Assignment	MUST NOT	There is no channel bonding in EPON.
58	SW Upgrade IPv6 TFTP Server	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
59	TFTP Server Provisioned Modem IPv6 Address	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
60	Upstream Drop Packet Classification	MUST	See the Upstream Drop Classification (TLV 60) section for more details
61	Subscriber Mgmt CPE IPv6 Prefix List	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
62	Upstream Drop Classifier Group ID	MUST	
63	Subscriber Mgmt Control Max CPE IPv6 Prefix	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
64	CMTS Static Multicast Session Encoding	SHOULD NOT	Multicast is not supported in this version of the specifications.
65	L2VPN MAC Aging Encoding	MUST	
66	Management Event Control Encoding	SHOULD NOT	Not needed for this version of the specifications.
79	UNI Control Encodings	MUST	
255	End-of-Data	MUST	

C.4 TLV 11³⁸

The following table specifies those TLV-11 CM SNMP MIB objects that MUST be supported by the DPoE System. The intent is to list only those objects that are currently in use (or planned to be in use) in CM provisioning files.

TLV-11 OID Name	Support Needed	Comments
docsDevFilterIpDefault	MUST	A value of 'accept' MUST be supported. A value of 'discard' MAY be supported. A DPoE System MUST ignore TLV 11 with the docsDevFilterIpDefault SNMP object with the value of 'discard', if this value is not supported.
docsDevFilterIpStatus	MUST	
docsDevFilterIpControl	MUST	'Policy' control is not required. Only 'accept' and 'discard' MUST be supported in this version of the specifications.
docsDevFilterIpPfIndex	MUST	
docsDevFilterIpDirection	MUST	
docsDevFilterIpBroadcast	MUST	A value of 'false' MUST be supported. A value of 'true' MAY be supported.
docsDevFilterIpSaddr	MUST	
docsDevFilterIpSmask	MUST	
docsDevFilterIpDaddr	MUST	

³⁸ Revised per MULPIv1.0-N-14.0149-1 on 5/27/14 by JB.

TLV-11 OID Name	Support Needed	Comments
docsDevFilterIpDmask	MUST	
docsDevFilterIpProtocol	MUST	
docsDevFilterIpSourcePortLow	MUST	
docsDevFilterIpSourcePortHigh	MUST	
docsDevFilterIpDestPortLow	MUST	
docsDevFilterIpDestPortHigh	MUST	
docsDevFilterIpTos	MUST	
docsDevFilterIpTosMask	MUST	
docsDevFilterIpContinue	SHOULD NOT	No "policy" treatment.
docsDevFilterIpPolicyId	SHOULD NOT	No "policy" treatment.
docsDevFilterLLCUnmatchedAction	MUST	
docsDevFilterLLCIfIndex	MUST	
docsDevFilterLLCProtocol	MUST	
docsDevFilterLLCProtocolType	MUST	
docsDevFilterLLCStatus	MUST	
docsDevNmAccessCommunity	MUST	
docsDevNmAccessControl	MUST	
docsDevNmAccessInterfaces	MUST	
docsDevNmAccessIp	MUST	
docsDevNmAccessIpMask	MUST	
docsDevNmAccessStatus	MUST	

C.5 Security (TLV 17)

TLV	Name	Support Needed	Comments
17.1	Authorize Wait Timeout	MUST NOT	None of these apply for EPON.
17.2	Reauthorize Wait Timeout	MUST NOT	None of these apply for EPON.
17.3	Authorization Grace Time	MUST NOT	None of these apply for EPON.
17.4	Operational Wait Timeout	MUST NOT	None of these apply for EPON.
17.5	Rekey Wait Timeout	MUST NOT	None of these apply for EPON.
17.6	TEK Grace Time	MUST NOT	None of these apply for EPON.
17.7	Authorize Reject Wait Timeout	MUST NOT	None of these apply for EPON.
17.8	SA Map Wait Timeout	MUST NOT	None of these apply for EPON.
17.9	SA Map Max Retries	MUST NOT	None of these apply for EPON.

C.6 Classification (TLVs 22 & 23)³⁹

TLV	Name	Support Needed	Comments
22.1 23.1	Classifier Reference	MUST	

³⁹ Revised per MULPIv1.0-N-12.0028-2 on 5/29/12 and per MULPIv1.0-N-12.0046-1 on 10/31/12 by JB.

TLV	Name	Support Needed	Comments
22.3 23.3	SF Reference	MUST	
22.5 23.5	Rule Priority	MUST	
22.6 23.6	Classifier Activation State	SHOULD NOT	Without deferred activation, there's no reason to support this TLV in this version of the specifications.
22.9.1 23.9.1	IPv4 TOS Range and Mask	MUST	
22.9.2 23.9.2	IP Protocol	MUST	
22.9.3 23.9.3	IPv4 Source Address	MUST	
22.9.4 23.9.4	IPv4 Source Mask	MUST	Mask can't be arbitrary bits. Must be used to define "most significant bits" a la IPv6 Prefix.
22.9.5 23.9.5	IPv4 Destination Address	MUST	
22.9.6 23.9.6	IPv4 Destination Mask	MUST	Mask can't be arbitrary bits. Must be used to define "most significant bits" a la IPv6 Prefix.
22.9.7 23.9.7	TCP/UDP Source Port Start	MUST	
22.9.8 23.9.8	TCP/UDP Source Port End	MUST	
22.9.9 23.9.9	TCP/UDP Destination Port Start	MUST	
22.9.10 23.9.10	TCP/UDP Destination Port End	MUST	
22.10.1 23.10.1	Ethernet DMAC	MUST	Mask can't be arbitrary bits. Must be used to define "most significant bits" a la IPv6 Prefix.
22.10.2 23.10.2	Ethernet SMAC	MUST	
22.10.3 23.10.3	Ethertype/DSAP/MacType	MUST	type=1 MUST be supported. type=3 does not apply to EPON. Other values MAY be supported.
22.11.1 23.11.1	[802.1p] User Priority	MUST	
22.11.2 23.11.2	[802.1q] VLAN ID	MUST	
22.12.1 23.12.1	IPv6 Traffic Class Range and Mask	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
22.12.2 23.12.2	IPv6 Flow Label	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
22.12.3 23.12.3	IPv6 Next Header Type	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
22.12.4 23.12.4	IPv6 Source Address	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
22.12.5 23.12.5	IPv6 Source Prefix Length	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
22.12.6 23.12.6	IPv6 Destination Address	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
22.12.7 23.12.7	IPv6 Destination Prefix Length	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
22.13	CM Interface Mask (CMIM) Encoding	MUST	
22.14 23.14	[802.1ad] S-VLAN Packet Classification Encodings	MUST	TLV for [802.1ad] S-VLAN classifier.

TLV	Name	Support Needed	Comments
22.14.1 23.14.1	[802.1ad] S-VLAN TPID	MUST	
22.14.2 23.14.2	[802.1ad] S-VLAN VID	MUST	
22.14.5 23.14.5	[802.1ad] C-VLAN TPID	MUST	
22.14.6 23.14.6	[802.1ad] C-VLAN VID	MUST	
22.15 23.15	[802.1ah] I-TAG Packet Classification Encodings	MUST	TLV for [802.1ah] I-TAG classifier.
22.15.1 23.15.1	[802.1ah] I-TAG I-TPID	MUST	
22.15.2 23.15.2	[802.1ah] I-TAG I-SID	MUST	
22.43 23.43	Vendor Specific Classifier Parameters	MUST	Type 43 TLVs within classifiers MUST be supported for purposes of MEF configuration. See [DPoE-SP-MEFv1.0].
23.43.5.1	VPN Identifier	MUST	
23.43.8	General Extension Information	MUST	Vendor ID of 0xFFFFFFFF

C.7 Service Flows (TLVs 24 & 25)⁴⁰

TLV	Name	Support Needed	Comments
24.1 25.1	SF Reference	MUST	
24.4 25.4	Service Class Name	MUST	
24.6 25.6	Quality of Service Parameter Set Type	MUST	A value of 0x7 MUST be supported. Registration MUST be denied if a value other than 0x7 is received.
24.7 25.7	Traffic Priority	MUST	
24.8	Upstream Max Sustained Traffic Rate	MUST	
25.8	Downstream Max Sustained Traffic Rate	MUST	
24.9 25.9	Maximum Traffic Burst	MUST	
24.10 25.10	Minimum Reserved Traffic Rate	MUST	
24.11 25.11	Assumed Min Rate Packet Size	SHOULD NOT	
24.12 25.12	Timeout for Active QoS Parameters	SHOULD NOT	There is no deferred admission/activation behavior in this version of the specifications.
24.13 25.13	Timeout for Admitted QoS Parameters	SHOULD NOT	There is no deferred admission/activation behavior in this version of the specifications.
24.14	Maximum Concatenated Burst	MUST NOT	This is a DOCSIS network-specific parameter which does not translate to a DPoE network.
25.14	Maximum Downstream Latency	SHOULD NOT	

⁴⁰ Revised per MULPIv1.0-N-12.0012-1 on 3/27/12, per MULPIv1.0-N-12.0028-2 on 5/29/12, and per MULPIv1.0-N-13.0105-1 on 10/4/13 by JB. Revised per MULPIv1.0-N-14.0117-2 on 2/18/14 and per MULPIv1.0-N-14.0127-1 on 2/27/14 by JB.

TLV	Name	Support Needed	Comments
24.15	SF Scheduling Type	MUST	Two values MUST be supported for this version of the specifications: 2 for Best Effort 4 for Real Time Polling Service Other values MAY be supported.
24.16	Request/Transmission Policy	MUST	The only bit that MUST be supported is Bit 4 (must not piggyback requests).
24.17	Nominal Polling Interval	MUST	This TLV is used with the Real Time Polling Scheduling Type.
25.17	Downstream Resequencing	MUST NOT	Not applicable to EPON.
24.18	Tolerated Poll Jitter	SHOULD NOT	
24.19	Unsolicited Grant Size	SHOULD NOT	This TLV is not needed to support the BE and RTP scheduling types.
24.20	Nominal Grant Interval	SHOULD NOT	This TLV is not needed to support the BE and RTP scheduling types.
24.21	Tolerated Grant Jitter	SHOULD NOT	This TLV is not needed to support the BE and RTP scheduling types.
24.22	Grants Per Interval	SHOULD NOT	This TLV is not needed to support the BE and RTP scheduling types.
24.23 25.23	IP ToS Overwrite	MUST	A "tos-and-mask" value of 0x00 MUST be supported. Other values MAY be supported.
24.26	Multiplier to Number of Bytes Requested	MUST NOT	Not applicable to EPON.
24.27	Upstream Peak Traffic Rate	SHOULD NOT	
25.27	Downstream Peak Traffic Rate	SHOULD NOT	
24.31 25.31	SF Required Attribute Mask	MUST	This attribute is used to steer SFs to an IP-SG.
24.32 25.32	SF Forbidden Attribute Mask	MUST NOT	The use of this attribute is not required, though may be used to, steer SFs to an IP-SG.
24.33 25.33	SF Attribute Aggregation Rule Mask	MUST NOT	Not applicable to EPON.
24.34 25.34	Application Identifier	SHOULD NOT	This TLV is used to extend admission control decisions based on PCMM policies. There is no PCMM support in this version of the specifications.
24.41 25.41	Data Rate Unit Setting	MUST	This setting provides a new data rate multiplier TLV as follows: bits/s → Kbits/s → Mbit/s → Gbit/s, increasing the maximum data rate to the maximum value of 232-1 Gbit/s.
24.43	Vendor Specific QoS Parameters	MUST	Type 43 TLVs within service flows MUST be supported for purposes of MEF configuration. See [DPoE-SP-MEFv1.0].
24.43.5.1	VPN Identifier	MUST	
24.43.8	General Extension Information	MUST	Vendor ID of 0xFFFFF.

C.8 Device Management (TLVs 38, 53 & 54)⁴¹

TLV	Name	Support Needed	Comments
34.1	SNMPv3 Kickstart Security Name	SHOULD	
34.2	SNMPv3 Kickstart Manager Public Number	SHOULD	
38.1	SNMPv3 Notification Receiver IPv4 Address	MUST	

⁴¹ Revised per MULPIv1.0-N-12.0024-1 on 5/29/12 by JB.

TLV	Name	Support Needed	Comments
38.2	SNMPv3 Notification Receiver UDP Port Number	MUST	
38.3	SNMPv3 Notification Receiver Trap Type	MUST	
38.4	SNMPv3 Notification Receiver Timeout	MUST	
38.5	SNMPv3 Notification Receiver Retries	MUST	
38.6	SNMPv3 Notification Receiver Filtering Parameters	MUST	
38.7	SNMPv3 Notification Receiver Security Name	MUST	
38.8	SNMPv3 Notification Receiver IPv6 Address	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
53.1	SNMPv1v2c Community Name	MUST	
53.2	SNMPv1v2c Transport Address Access	MUST	
53.2.1	SNMPv1v2c Transport Address	MUST	
53.2.2	SNMPv1v2c Transport Address Mask	MUST	
53.3	SNMPv1v2c Access View Type	MUST	
53.4	SNMPv1v2c Access View Name	MUST	
54.1	SNMPv3 Access View Name	MUST	
54.2	SNMPv3 Access View Subtree	MUST	
54.3	SNMPv3 Access View Mask	MUST	
54.4	SNMPv3 Access View Type	MUST	

C.9 TLV 43⁴²

TLV	Name	Support Needed	Comments
43.1	CM Load Balancing Policy ID	MUST NOT	Not applicable to EPON.
43.2	CM Load Balancing Priority	MUST NOT	Not applicable to EPON.
43.3	CM Load Balancing Group ID	MUST NOT	Not applicable to EPON.
43.4	CM Ranging Class ID Extension	MUST NOT	Not applicable to EPON.
43.5	L2VPN Encoding	MUST	
43.6	Extended CMTS MIC Configuration Setting	MAY	ONUs do not request config files, thus tampering is not a concern. This TLV may be supported for backwards compatibility.
43.7	Source Address Verification Authorization Encoding	MUST	
43.8	General Extension Information	MUST	
43.9	Cable Modem Attribute Masks	MUST NOT	This is really bonding-specific, thus not applicable to EPON.
43.10	IP Multicast Join Authorization Encoding	SHOULD NOT	Multicast is not supported in this version of the specifications.
43.11	Service Type Identifier	SHOULD NOT	This TLV is not applicable to DPoE Networks in this version of the specifications.

⁴² Revised per MULPIv1.0-N-12.0011-1 on 3/27/12 by JB.

C.10 [DPoE-SP-MEFv1.0] and [L2VPN] (TLVs 43.5, 45, & 65)

See [DPoE-SP-MEFv1.0] for more information on the expected use and support for these TLVs in this version of DPoE specifications.

TLV	Name	Support Needed	Comments
43.5.1	VPN Identifier	MUST	
43.5.2	NSI Encapsulation Subtype	MUST	
43.5.2.1	Other	SHOULD NOT	
43.5.2.2	[802.1q]	MUST	
43.5.2.3	[802.1ad]	MUST	
43.5.2.4	MPLS Peer	SHOULD NOT	
43.5.2.5	L2TPv3 Peer	SHOULD NOT	
43.5.2.6	[802.1ah]	MUST	
43.5.2.6.1	[802.1ah] I-Tag	MUST	
43.5.2.6.2	[802.1ah] B-DA	MUST	
43.5.2.6.3	[802.1ah] B-SA	SHOULD NOT	
43.5.3	Enable eSAFE DHCP Snooping	MUST	
43.5.4	CM Interface Mask	MUST	
43.5.5	Attachment Group ID	SHOULD NOT	
43.5.6	Source Attachment Individual ID	SHOULD NOT	
43.5.7	Target Attachment Individual ID	SHOULD NOT	
43.5.8	Ingress User Priority	SHOULD NOT	
43.5.9	User Priority Range	SHOULD NOT	
43.5.10	L2VPN SA-Descriptor	MUST NOT	BPI does not apply here.
43.5.13	L2VPN Mode	MUST	
43.5.14	DPoE TPID Translation	MUST	
43.5.14.1	Upstream TPID	MUST	
43.5.14.2	DownstreamTPID	MUST	
45.1	Downstream Unencrypted Traffic (DUT) Control	SHOULD NOT	
45.2	Downstream Unencrypted Traffic (DUT) CMIM	SHOULD NOT	
65.1	L2VPN MAC Aging Mode	MUST	

C.11 Customer (Subscriber) Management (TLV 43.7)

TLV	Name	Support Needed	Discussion
43.7.1	SAV Group Name Subtype	MUST	
43.7.2	SAV Static Prefix Rule Subtype	MUST	
43.7.2.1	SAV Static Prefix Address Subtype	MUST	IPv4 mode only.
43.7.2.2	SAV Static Prefix Length Subtype	MUST	IPv4 mode only.

C.12 Upstream Drop Classification (TLV 60)

TLV	Name	Support Needed	Comments
60.1	Classifier Reference	MUST	
60.5	Rule Priority	MUST	

TLV	Name	Support Needed	Comments
60.9.1	IPv4 TOS Range and Mask	MUST	
60.9.2	IP Protocol	MUST	
60.9.3	IPv4 Source Address	MUST	
60.9.4	IPv4 Source Mask	MUST	Mask can't be arbitrary bits. Must be used to define "most significant bits" a la IPv6 Prefix.
60.9.5	IPv4 Destination Address	MUST	
60.9.6	IPv4 Destination Mask	MUST	Mask can't be arbitrary bits. Must be used to define "most significant bits" a la IPv6 Prefix.
60.9.7	TCP/UDP Source Port Start	MUST	
60.9.8	TCP/UDP Source Port End	MUST	
60.9.9	TCP/UDP Destination Port Start	MUST	
60.9.10	TCP/UDP Destination Port End	MUST	
60.10.1	Ethernet DMAC	MUST	Mask can't be arbitrary bits. Must be used to define "most significant bits" a la IPv6 Prefix.
60.10.2	Ethernet SMAC	MUST	
60.10.3	Ethertype/DSAP/MacType	MUST	type=1 must be supported. type=3 does not apply to EPON. Other type values MAY be supported.
60.11.1	[802.1p] User Priority	MUST	
60.11.2	[802.1q] VLAN ID	MUST	
60.12.1	IPv6 Traffic Class Range and Mask	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
60.12.2	IPv6 Flow Label	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
60.12.3	IPv6 Next Header Type	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
60.12.4	IPv6 Source Address	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
60.12.5	IPv6 Source Prefix Length	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
60.12.6	IPv6 Destination Address	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
60.12.7	IPv6 Destination Prefix Length	SHOULD NOT	There is no support for IPv6 in this version of the specifications.
60.13	CM Interface Mask (CMIM) Encoding	MUST	
60.43	Vendor Specific Classifier Parameters	MUST	These TLVs will be "supported" in that vendor-proprietary TLVs can be specified in the configuration file, but their implementation is up to the vendor. There currently are no standard DPoE vendor-specific classifier parameters.

Annex D eSAFE DHCP Snooping⁴³

This section specifies the requirements around enabling eSAFE DHCP snooping as defined in [L2VPN].

- TLV 43.5.3 "eSAFE DHCP Snooping" defined in [L2VPN] MUST be supported by all DPoE System Relay Agents. DPoE Standalone ONUs with one or more eSAFE's present configured (on) with TLV 43.5.3 MUST relay eSAFE DHCP messages from the DPoE Standalone ONU to the DPoE System on an EVC with S-VLAN ID and C-VLAN ID as configured for the corresponding CMIM or S₁ interface.

⁴³ Revised per MULPIv1.0-N-12.0017-2 on 5/29/12 by JB.

Annex E MPCP Discovery Processing in DPoE Networks⁴⁴

E.1 IEEE MPCP Discovery Process

The ONU Discovery Process for EPON defined in [802.3] involves the exchange of multiple messages between OLT and ONU. A diagram showing the message sequence of the MPCP Discovery Process is provided in Figure 9.

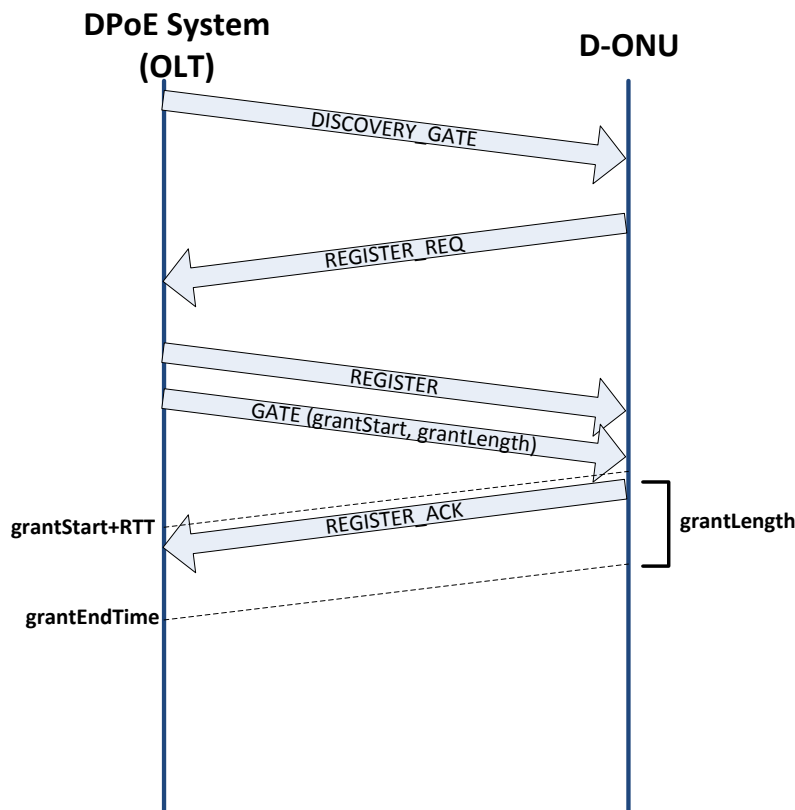


Figure 9 - MPCP Discovery Process Message Sequence.

The Discovery Process begins with the OLT sending a Discovery GATE message to open a contention window for ONUs to transmit a REGISTER_REQ message. Upon receipt of the REGISTER_REQ message, the OLT responds with a REGISTER message. Shortly after sending the REGISTER message, the OLT sends a GATE message to the registering ONU. The ONU uses the corresponding grant in the GATE message to transmit the REGISTER_ACK message. If OLT receives the REGISTER_ACK message by *grantEndTime* then the ONU is registered and the Discovery Process is complete. If the REGISTER_ACK message is not received by *grantEndTime*, the ONU is de-registered and must restart the Discovery Process by waiting for a Discovery GATE message.

The time by which the OLT must receive the REGISTER_ACK message is indicated in the variable *grantEndTime*. While the [802.3] standard defines how to compute *grantEndTime* based on when the GATE message was sent, the standard does not indicate how long the OLT should wait after sending the REGISTER message before it sends the GATE message. It is recognized that ONUs require a variable amount of processing time before it can respond with the REGISTER_ACK message. This variable and non-standardized processing time leads to interoperability challenges, which motivates changes to the Discovery Process in DPoE Networks, as outlined below.

⁴⁴ Added Annex E per MULPIv1.0-N-13.0079-1 on 6/19/13 by JB.

E.2 DPoE MPCP Discovery Process

The DPoE MPCP Discovery Process implements several modifications aimed at increasing the robustness of the D-ONU discovery. The DPoE MPCP Discovery Process specifies a delay that the DPoE system must enforce between issuing the REGISTER message to a registering D-ONU and issuing a GATE message to the same D-ONU to provide it an opportunity to transmit the REGISTER_ACK message.

To demonstrate the modifications required to the Discovery Process in DPoE Networks, several new variables are defined:

- *registerStartTime* – the time on DPoE System when REGISTER message is sent to a D-ONU
- *registerAckDelay* – a delay after which DPoE System sends a GATE message for D-ONU to respond with REGISTER_ACK message.
- *registerDeadline* – the time on DPoE System indicating the registration deadline. If the REGISTER_ACK message is not received by this time, the DPoE System deregisters the D-ONU.

The DPoE MPCP Discovery Process specifies a delay (*registerAckDelay*) that the DPoE system must enforce between issuing the REGISTER message to a registering D-ONU (*registerStartTime*) and issuing a GATE message to the same D-ONU to provide it an opportunity to transmit the REGISTER_ACK message. The first GATE message transmitted at or after DPoE System time (*registerStartTime+registerAckDelay*) is denoted the “Deadline GATE”. The *registerDeadline* value is set to be equal to the *grantEndTime* of the first grant in the Deadline GATE message.

These time variables are related to each other in the following manner:

$$DeadlineGATE.timestamp \geq registerStartTime + registerAckDelay$$

$$registerDeadline = DeadlineGATE.grantStart[0] + DeadlineGATE.grantLength[0] + RTT$$

The factor *RTT* is the round trip time to the D-ONU. The relationship between these time variables is shown pictorially in Figure 10.

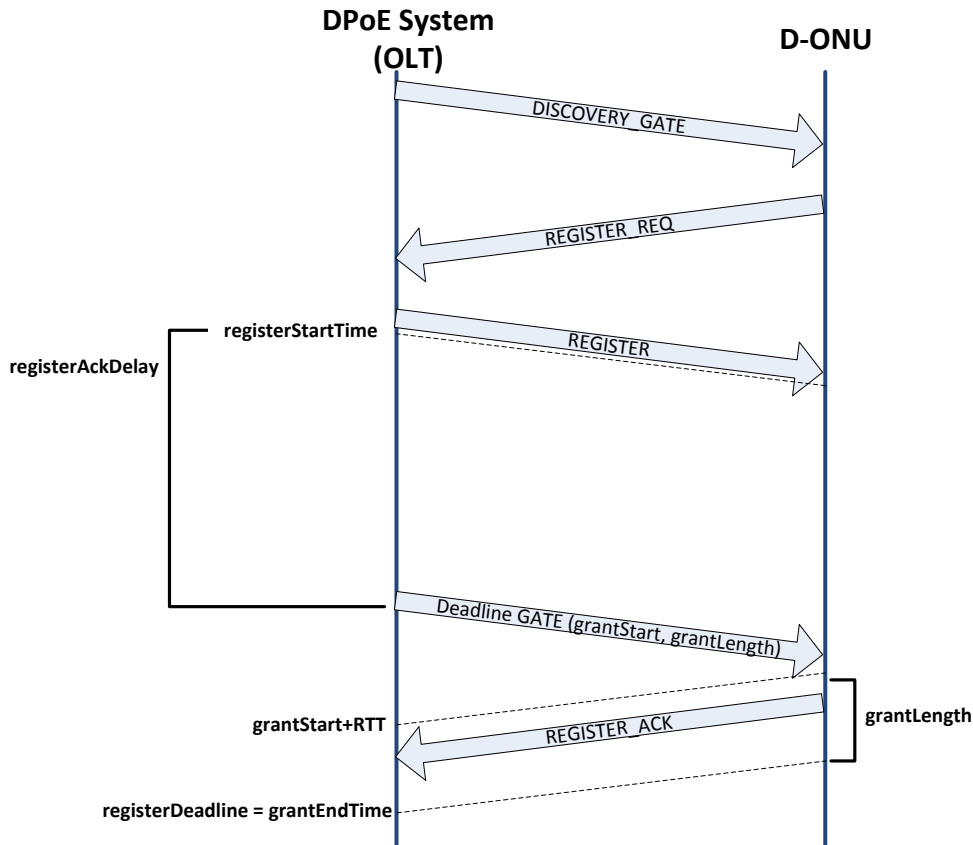


Figure 10 - Discovery Process with RegisterStartTime, RegisterAckDelay, and RegisterDeadline.

The DPoE System MUST employ a mechanism that allows for D-ONUs at least *registerAckDelay* time to respond with the REGISTER_ACK message once the D-ONU has received the REGISTER message. The DPoE System MUST use a default value of 100 ms for *registerAckDelay*. The DPoE System MUST provide the capability to configure *registerAckDelay* to a value different from the default value. If DPoE System does not receive the REGISTER_ACK message by the *registerDeadline*, the DPoE System MUST deregister the D-ONU. The specific mechanism employed by the DPoE System to meet these requirements is vendor-specific.

The D-ONU MUST transmit the REGISTER_ACK message within 100 ms after receiving the REGISTER message.

E.2.1 Discovery Using Multiple GATE messages

If a D-ONU does not require the entire *registerAckDelay* interval to generate the REGISTER_ACK message, it is beneficial to complete the Discovery process of such D-ONU sooner. To allow such expedited discovery, the DPoE System may optionally issue one or more GATE messages before the Deadline GATE message. GATE messages transmitted prior to the Deadline GATE message are called Auxiliary GATE messages.

Figure 11 shows two Auxiliary GATE messages sent to the D-ONU. In this example, the D-ONU is unable to generate the REGISTER_ACK message quickly enough to respond in the grant contained within the first Auxiliary GATE message. Consequently, the DPoE System sends a second Auxiliary GATE message, which is used by the D-ONU to transmit the REGISTER_ACK message.

The number of Auxiliary GATE messages and the time period between GATE messages is vendor specific. The discovery process completes as soon as the REGISTER_ACK message is received by the DPoE System, which can occur significantly ahead of the *registerDeadline* time. If the REGISTER_ACK message is received in a grant in one of the Auxiliary GATE messages, the DPoE System does not need to continue issuing any additional Auxiliary GATE messages or the Deadline GATE message.

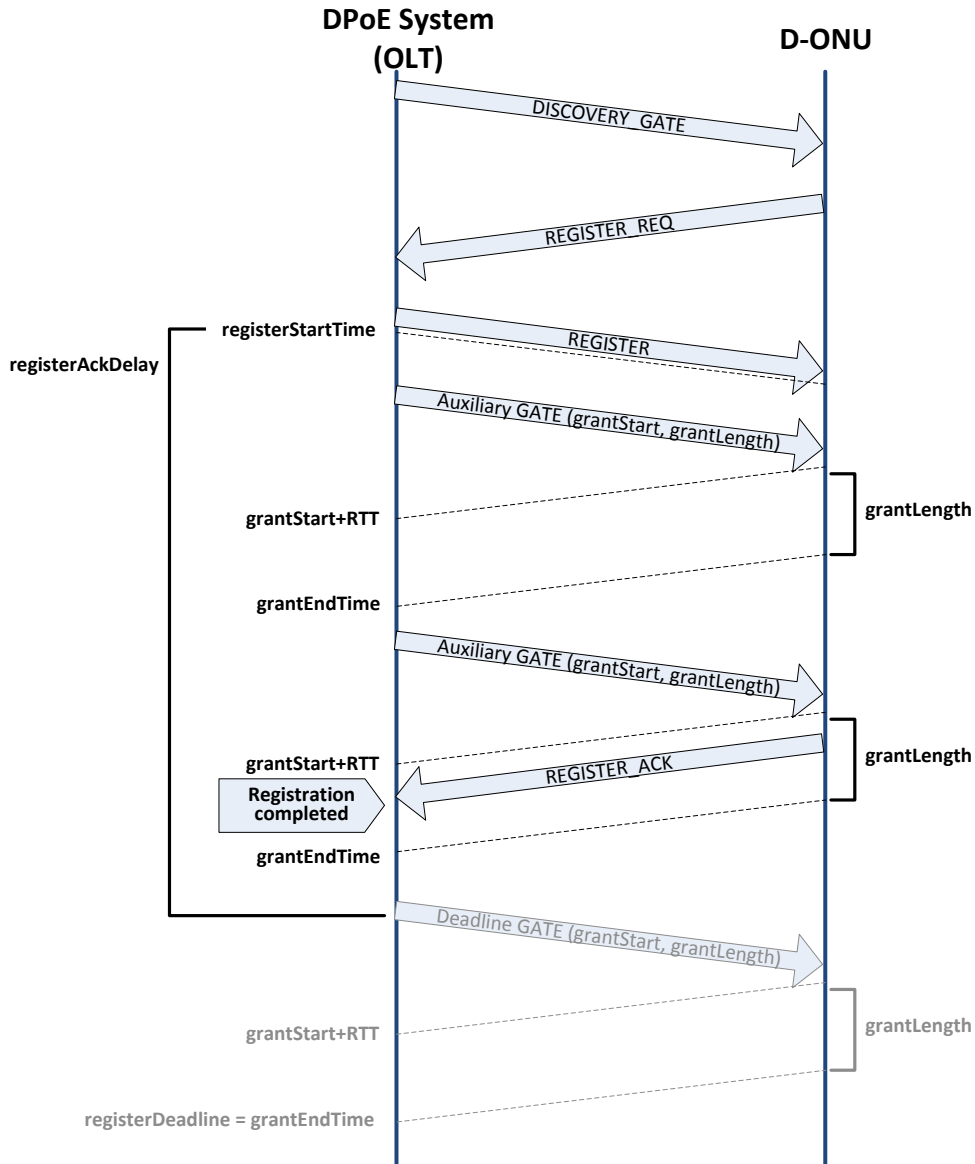


Figure 11 - DPoE MPCP Discovery Process Using Multiple Auxiliary GATE Messages, REGISTER_ACK Message Arrived in Grant from Auxiliary GATE Message.

Figure 12 shows a scenario where the D-ONU is unable to generate the REGISTER_ACK message quickly enough to respond in the grant contained within any of the Auxiliary GATE messages. After the expiration of the *registerAckDelay* interval, the DPoE system issues the Deadline GATE message, as described above. In this example, the REGISTER_ACK message is received during the grant allocated by the Deadline GATE message.

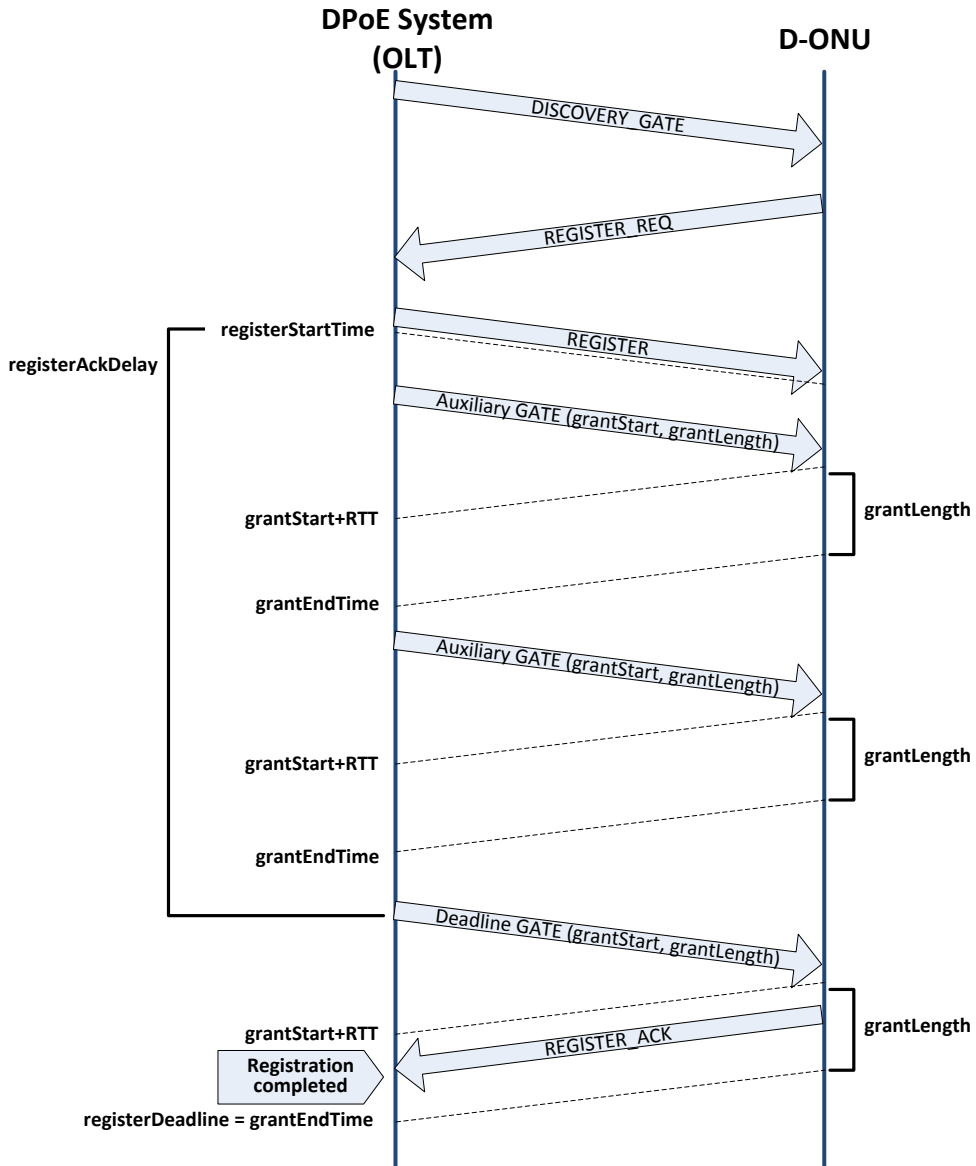


Figure 12 - DPoE MPCP Discovery Process using multiple GATE messages; REGISTER_ACK message arrived in grant from Deadline GATE message.

Appendix I EPON Media Access Control Protocol (Informative)⁴⁵

The DPoE specifications rely on the respective [802.3] and [802.3] standards and do not introduce any changes to their stipulations.

I.1 Timing and Synchronization

In an EPON network, timing and synchronization services for the EPON transport (over the TU interface) are provided by the Multi-Point Control Protocol (MPCP) as defined in [802.3ah], Clause 64 and [802.3], Clause 77.

The ranging and timing process in use in 1G-EPON is defined in [802.3] subclause 64.2.1.1, while the delay variation requirements are defined in [802.3] subclause 65.3.3. Similarly, [802.3] subclause 77.2.1.1 defines the ranging and timing process in use in 10G-EPON, with delay variation specified in [802.3] subclause 76.5.4.11.

I.1.1 MPCP Clock Synchronization

In order to decouple the time of GATE MPCPDU transmission from the start time of the granted upstream slot, the OLT and each ONU maintain a local clock, commonly referred to as the *MPCP clock*. This MPCP clock has the form of a 32-bit counter, increasing every time quanta (TQ), defined to be equal to 16 ns, or the time required to transmit 2 bytes of data at data rate of 1 Gbps. Effectively, a number of EPON parameters, including the timeslot start times and lengths carried in GATE MPCPDUs, as well as queue lengths carried in REPORT MPCPDUs, are expressed in the units of TQ.

Synchronization between the ONU MPCP clock and the OLT MPCP clock is carried out through the exchange of MPCPDUs. In this process, the OLT MPCP clock is considered to be the timing master and the ONU MPCP clock is always slaved to the OLT time domain. Each MPCPDU carries a timestamp of the local MPCP clock at the time when the given MPCPDU passes through the MAC Control Multiplexer. For details of MPCPDU structure and functional specifications, see [802.3ah] subclause 64.3.6 and [802.3] subclause 77.3.6.

This clock synchronization scheme is based on the assumption that frame propagation delay between the Control Multiplexer at the transmitting device and the Control Parser at the receiving device is nearly constant. In other words, frames cannot be blocked or delayed in the MAC and PHY sub-layers.

I.1.2 Loop Timing in EPON

In traditional P2P Ethernet, local oscillators are allowed to deviate from the nominal frequency by 100 ppm. Such a relaxed clock tolerance allowed for very inexpensive devices to be built. In EPON, to remedy this situation, MPCP mandates strict loop timing for the ONU, which means that MPCP clocks operating in all slaved ONUs must track the received OLT MPCP clock, recovered from the data transmitted by the OLT. Under such a scenario, even though the OLT MPCP clock is still allowed to deviate ± 100 ppm from the nominal frequency, ONUs are able to recover the clock and remain synchronized at all times. This is achieved thanks to the fact that the downstream channel OLT constantly transmits data or idle characters, meaning that ONUs never lose synchronization OLT clock as long as they are able to receive and recover data.

⁴⁵ Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

Appendix II EPON Multipoint Control Protocol Data Units (Informative)⁴⁶

The Multi-Point Control Protocol (MPCP) was specified in [802.3] to resolve the problems related with P2P Ethernet operation in the P2MP environment of EPON systems. MPCP is used in EPON to dynamically allocate access to the transmission medium (ODN path) to individual ONUs connected to EPON. It assigns upstream transmission slots to all active slave devices. Provided that stable operation conditions are maintained in the network and no link suffers from significant variations of the Round Trip Time (RTT), the allocated slots are always non-overlapping. This means that, upon their arrival at the OLT receiver module, the data frames can be received, delineated, and decoded. MPCP provides the complete signaling infrastructure (control plane) for coordinating data transmissions originating from ONUs to an OLT. The functionality of the MPCP sublayer in the ONU and the OLT is quite different:

- The OLT MPCP sublayer is responsible for Discovery of the newly connected stations, their registration, and measurement of the RTT, as well as scheduling and controlling the transmission from individual ONUs in the upstream channel;
- The ONU MPCP sublayer is mainly responsible for the reporting of the current queue state at the end of the upstream transmission slot (provided that the OLT MPCP requested such functionality through the respective GATE MPCPDU), as well as participation in the Discovery process.

The principle of operation for the MPCP mechanism is relatively straightforward. The total available upstream channel bandwidth is divided into transmission units (typically termed *slots*) using the TDMA technique. Each such slot of an arbitrary length can be assigned to an ONU (more specifically to the respective LLIDs) based on the DBA mechanism under operation in the OLT central packet scheduler. The scheduler assigns each LLID a certain fraction of the upstream transmission slot, which depends on the current bandwidth demand of the given entity (as indicated using the REPORT MPCPDU), available bandwidth, bandwidth demand of other LLIDs, number of LLIDs, employed service policy, etc. The ONU is then notified of the size and start of the transmission slot using the complementary GATE MPCPDU.

The MPCP transmission arbitration is based on two messages, namely REPORT and GATE MPCPDUs. The REPORT MPCPDUs are transmitted by the ONU and are used to indicate the current bandwidth demand to the OLT. The bandwidth demand is typically estimated based on the current queue occupancy (a single ONU can hold a number of packet queues storing Ethernet frames, mapped into a number of available LLID entities [802.3]), with the maximum number of queue reports included in a single REPORT MPCPDU limited to 13 due to the finite and pre-defined size of an MPCPDU. EPON specifications allow additionally for existence of queue thresholds. These provide the ONU with the ability to indicate several delineation boundaries per single queue, increasing the scheduling efficiency at the OLT side by providing additional information on the internal structure of each particular queue. Queue threshold use and treatment are not specified in [802.3] and are left to vendors to optionally implement.

Once received at the OLT, the REPORT MPCPDU is parsed and passed to the DBA module responsible for scheduling the size and start time for upstream transmission slots. Scheduling must account for both the burst-mode delays and path delay variations between near and distant ONUs to prevent overlap of upstream transmissions at the OLT receiver. The size of each allocated slot depends on the actual bandwidth demand, selected service policy (whether static or dynamic bandwidth allocation is used), number of active LLIDs, amount of available bandwidth, polling protocol in use, etc. MPCP was designed to operate with any DBA algorithm. By requiring a common control plane, but allowing for different algorithms, any vendor can develop new bandwidth allocation protocols with arbitrary complexity. Once the DBA module completes the slot size and time estimation process, a GATE MPCPDU is constructed, loaded with the respective DBA-estimated information, and delivered downstream at the first possible opportunity. All MPCPDUs are transmitted with the highest priority, but may be queued after a long frame under transmission.

In accordance with the EPON specifications, a GATE MPCPDU allows the central DPoE System controller to schedule at most 4 transmission slots at once (so-called scheduling into the future), with the size of $2^{16}-1$ TQ (1 TQ = 2 B = 16 ns for effective 1 Gbps data rate), resulting in a single transmission slot limited to roughly 128 kB. Upon

⁴⁶ Revised per MULPIv1.0-N-14.0155-1 on 7/14/14 by JB.

reception of such an MPCPDU, the ONU updates its local clock index using the time-stamp field carried in the message body, thus effectively maintaining global synchronization with the OLT clock without the need for a separate clock signal. The scheduling information is parsed and processed accordingly, resulting in the creation of transmission events, which are executed once the local clock value reaches the slot start value, as indicated in the previously processed GATE MPCPDU. During a transmission slot, the given ONU delivers backlogged Ethernet frames using its local intra-ONU scheduler, attempting to fill in the allocated slot as much as possible. Since Ethernet frames cannot be fragmented and delineation bounds typically change between the REPORT MPCPDU transmission and reception of the respective GATE MPCPDU, unused slot remainders are created, leading to certain inefficiencies in the upstream channel transmission. The remaining frames, which do not fit the currently allocated slot, are delayed to the next transmission is granted by the OLT scheduler.

Appendix III Acknowledgements

On behalf of our industry, we would like to thank the following individuals for their contributions to the development of this specification.

Contributor	Company Affiliation
John Dickinson	Bright House Networks
Andrew Chagnon, Drew Davis, James Fletcher, Paul Gray	Broadcom
Mike Holmes	Broadway Networks
Curtis Knittle, Bob Lund, Glenn Russell	CableLabs
Shamim Akhtar, Doug Jones, Mehmet Toy	Comcast
Vladimir Bronstein, James Chen, Hesham ElBakoury	Hitachi Communication Technologies America
Victor Blake	Independent Consultant
Robert Harris, Karen Rice, Ashish Sardesai	Time Warner Cable
Marek Hajduczenia, Nevin Jones, Zang Meiyuan	ZTE

In addition to the above, we would like to thank the following individuals who provided subject matter expertise on one or more aspects of the project:

Contributor	Company Affiliation
Edwin Mallette	Bright House Networks
Andrew Dellow, Niki Pantelias, Ricki Li, Paul Runcy, Ed Boyd, Howard Abramson, Matt Hartling	Broadcom
Jianhui Zhou, Wen Li, Fulin Pan	Broadway Networks
Brian Hedstrom, Stuart Hoggan, Greg White, Karthik Sundaresan, Chris Donley	CableLabs
Jason Combs, Saif Rahman, Matt Scully, Bin Wen, Philip Chang, Rashid Siddiqui	Comcast
Dylan Ko, Simon Zhu, Guru Yeleswarapu, Jeff Stribling	Hitachi
Ron daSilva, Mike Kelsen, Shan Huang, Matt Cannon, Tushar Nakhre	Time Warner Cable
Stove Li Zhang, David Chen, Dick Chen	ZTE

Appendix IV Revision History

IV.1 Engineering Changes incorporated into DPOE-SP-MULPIv1.0-I02-120607

ECN	ECN Date	Summary	Author
MULPIv1.0-N-11.0005-1	5/05/2011	Missing classifier configuration C-VLAN TLV	Marek Hajduczenia
MULPIv1.0-N-12.0010-1	3/02/2012	Addition of Primary Service Flows for DPoE 1.0	Paul Gray
MULPIv1.0-N-12.0011-1	3/02/2012	Changing MIC TLV requirement	Karthik Sundaresan
MULPIv1.0-N-12.0012-1	3/15/2012	Removing requirement for Max Concatenation Burst	Karthik Sundaresan
MULPIv1.0-N-12.0017-2	4/26/2012	Correcting DHCPv4 Options	Karthik Sundaresan
MULPIv1.0-N-12.0022-1	5/03/2012	Removal of untestable / unnecessary requirements from DPoE-SP-MULPIv1.0	Marek Hajduczenia
MULPIv1.0-N-12.0024-2	5/03/2012	Clarification of L2Vpnlid for pairing of flows	Kirk Erichsen
MULPIv1.0-N-12.0028-2	5/10/2012	Clarification of L2Vpnlid for pairing of flows	Paul Gray
MULPIv1.0-N-12.0030-1	5/10/2012	Add two DPoE Requirements from DOCSIS MULPI	Stephen Schottler
MULPIv1.0-N-12.0034-1	5/10/2012	Informative description of LLID	Curtis Knittle

IV.2 Engineering Change incorporated into DPOE-SP-MULPIv1.0-I03-120830

ECN	ECN Date	Summary	Author
MULPIv1.0-N-12.0042-3	7/26/2012	Classification Depth for Layer 2 Classification	Karthik Sundaresan

IV.3 Engineering Change incorporated into DPOE-SP-MULPIv1.0-I04-121115

ECN	ECN Date	Summary	Author
MULPIv1.0-N-12.0046-1	10/18/2012	Clarification on the use of TLV [22/23]14.6 and [22/23]11.2	Marek Hajduczenia

IV.4 Engineering Change incorporated into DPOE-SP-MULPIv1.0-I05-130328

ECN	ECN Date	Summary	Author
MULPIv1.0-N-12.0058-1	1/17/2013	Changes to REQ15801 and REQ15802 in MULPIv1	Marek Hajduczenia

IV.5 Engineering Changes incorporated into DPOE-SP-MULPIv1.0-I06-130808

ECN	ECN Date	Summary	Author
MULPIv1.0-N-13.0079-1	5/23/2013	Discovery Timing in DPoE Networks	Curtis Knittle
MULPIv1.0-N-13.0083-1	5/30/13	TLV-29 Clarification	Stuart Hoggan

IV.6 Engineering Change incorporated into DPOE-SP-MULPIv1.0-I07-131114

ECN	ECN Date	Summary	Author
MULPIv1.0-N-13.0105-1	10/03/2013	EC to Support for SF Required Attribute Mask MULPIv1.0	Edwin Mallette

IV.7 Engineering Changes incorporated into DPOE-SP-MULPIv1.0-I08-140327

ECN	ECN Date	Summary	Author
MULPIv1.0-N-14.0117-2	02/06/2014	IP(HSD) Changes for MULPIv1.0	Paul Gray
MULPIv1.0-N-14.0125-2	02/27/2014	UNI management parameters in CM configuration file	Marek Hajduczenia
MULPIv1.0-N-14.0127-1	02/27/2014	Data Rate Unit Setting for DPoE 1.0	Lane Johnson
MULPIv1.0-N-14.0129-2	02/27/2014	CMIM clarification for DPoE.10	Lane Johnson
MULPIv1.0-N-14.0138-1	03/04/2014	DPoE UNI control Top TLV Update	Lane Johnson

IV.8 Engineering Changes incorporated into DPOE-SP-MULPIv1.0-I09-140807

ECN	ECN Date	Summary	Author
MULPIv1.0-N-14.0134-1	3/27/2014	Support for larger MTUs	Marek Hajduczenia
MULPIv1.0-N-14.0149-1	5/15/2014	docsDevFilterIpDefault and DPoE System behavior	Marek Hajduczenia
MULPIv1.0-N-14.0155-1	7/3/2014	Alignment and cleanup of 802.3 references	Marek Hajduczenia
