Cable Data Services
DOCSIS® Provisioning of EPON Specifications

DPoE™ Architecture Specification

DPoE-SP-ARCHv1.0-I01-110225

ISSUED

Notice

Document Status Sheet

Document Control Number:	DPoE-SP-ARCHv1.0-I01-110225			
Document Title:	DPoE™ Architecture Specification			
Revision History:	I01 - Released 2/25/11			
Date:	February 25, 2011			
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 stable document, which has undergone rigorous member and vendor review and is suitable for product design and development, cross-vendor interoperability, and for certification testing.

Closed A static document, reviewed, tested, validated, and closed to further engineering change requests to the specification through CableLabs.

Trademarks

Contents

Figures

Tables

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 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 Introduction

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 [802.3ah] and the amendment for 10G-EPON [802.3av] support a centralized operator-based controller (OLT) architecture 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 [802.3ah]. 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 an SID, supports both unicast and broadcast, and has other similarities.

Existing [802.3ah] EPON systems do interoperate within the strict definitions of 1G-EPON. Experience with lab testing, field trials, and deployments has shown operators that 1G-EPON OLT and ONU systems typically only interoperate with a single port ONU. This is because [802.3ah] specifies the interfaces on the PON (the DPoE TU interface) but does not specify any of the other system interfaces. For example, an OLT from vendor A will register an ONU from vendor B, but it is not possible to construct a VLAN from the DPoE MN interface to an S interface. This is a well-recognized limitation of [802.3ah]. The challenge is that neither 1G-EPON nor 10G-EPON specify OAMP to forward traffic between NNI ports and the PON, or UNI ports and the PON. This is not different from other Ethernet standards. For example, if two Ethernet switches from two different vendors are connected, each switch must typically be configured independently. The challenge for EPON is that the remote device (the ONU) cannot be reached, and therefore cannot be configured. A solution to this problem must then be based on developing a common (standard) method of reaching the controller for the ONU, identifying the ONU capabilities, and providing that information to the OLT so that it can configure a working end to forwarding service (in both directions).

Even if EPON had solved that provisioning challenge, there are no standard management interfaces for the ongoing operations and maintenance of the network, including fault management, performance management, security, etc. Operators already have fully working and scaled-out systems that solve these challenges for DOCSIS networks. One of the primary goals for DPoE specifications is to use the existing DOCSIS back office infrastructure to scale up EPON-based business services.

1.2 Scope

This specification describes the architecture for DPoE Networks and the organization of the DPoE specifications.

Existing business services include one or more of DOCSIS IP and Layer 2 services, baseband and broadband Ethernet services over coaxial cable, Ethernet over fiber with baseband and broadband (CWDM and DWDM), BPON, EPON, and wireless services. The majority of business services (and all residential Internet and voice) customers are supported by the DOCSIS systems and processes. The maturity of both the technology and the back office systems and process allows for a high degree of scaling as evidenced by the growth of IP(HSD) (residential broadband) and more recently voice service, using these existing processes and systems.

This version of the document specifies requirements for these services only:

- MEF EPL Service
- IP/Internet Service as defined by DOCSIS 3.0 (for IPv4 only)

Other services can be operated over the top of the above two services, but the provisioning, operations, and other requirements for such services are not specified in this version of DPoE. For example, an operator could use a DEMARC device or functions on a DPoE ONU outside of the scope of these specifications to manually or otherwise configure and operate other Metro Ethernet services, such as E-VPL, E-Tree, or E-LAN. This version of DPoE specifications does not specify the provisioning and operations of these services, but vendors could offer, or operators could implement, such "over the top" solutions.

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 (IP NE) Requirements [DPoE-SP-IPNEv1.0] 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 Specifications

This document is one in a series of eight (8) 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 [802.3ah] and [802.3av] 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.

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 are available through either IP or Layer 2 Ethernet-based Network Interfaces.

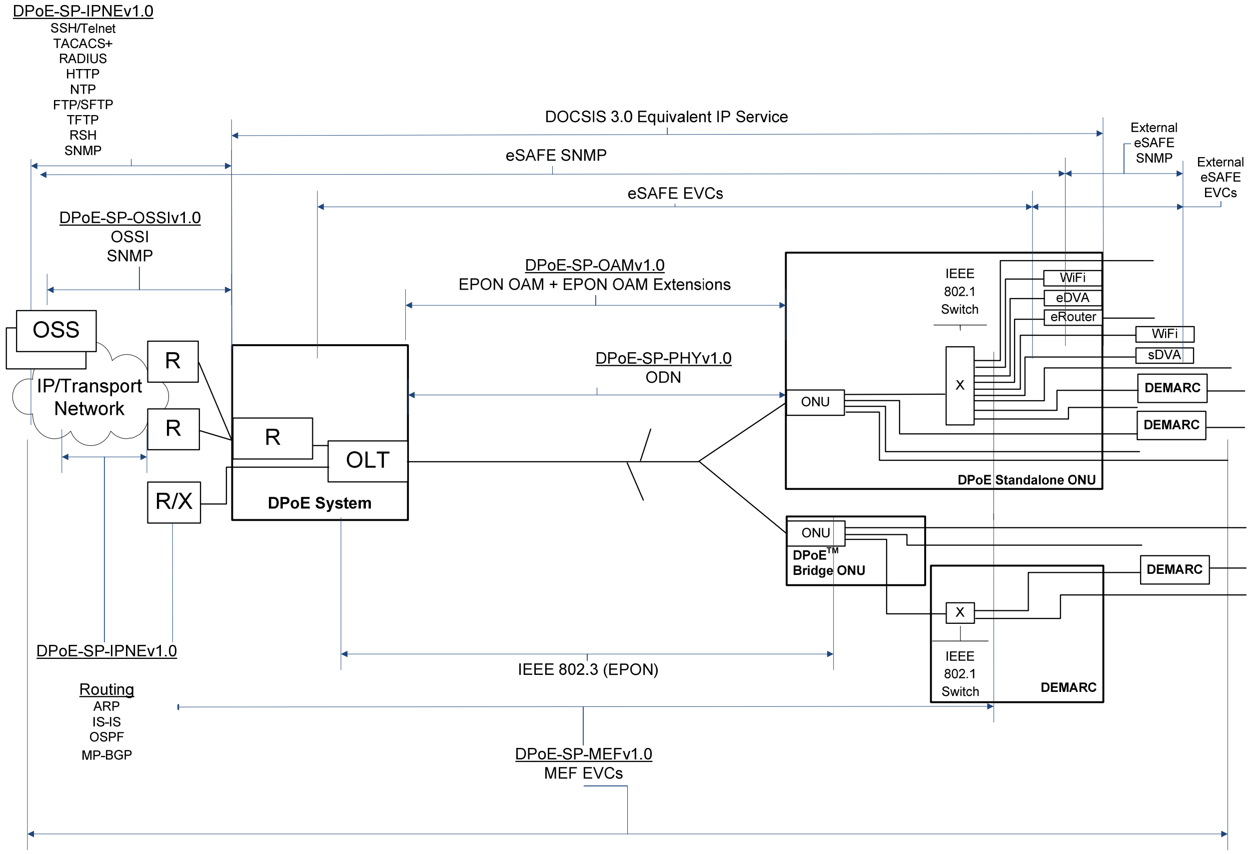

Figure 1 - DPoE Reference Architecture

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.

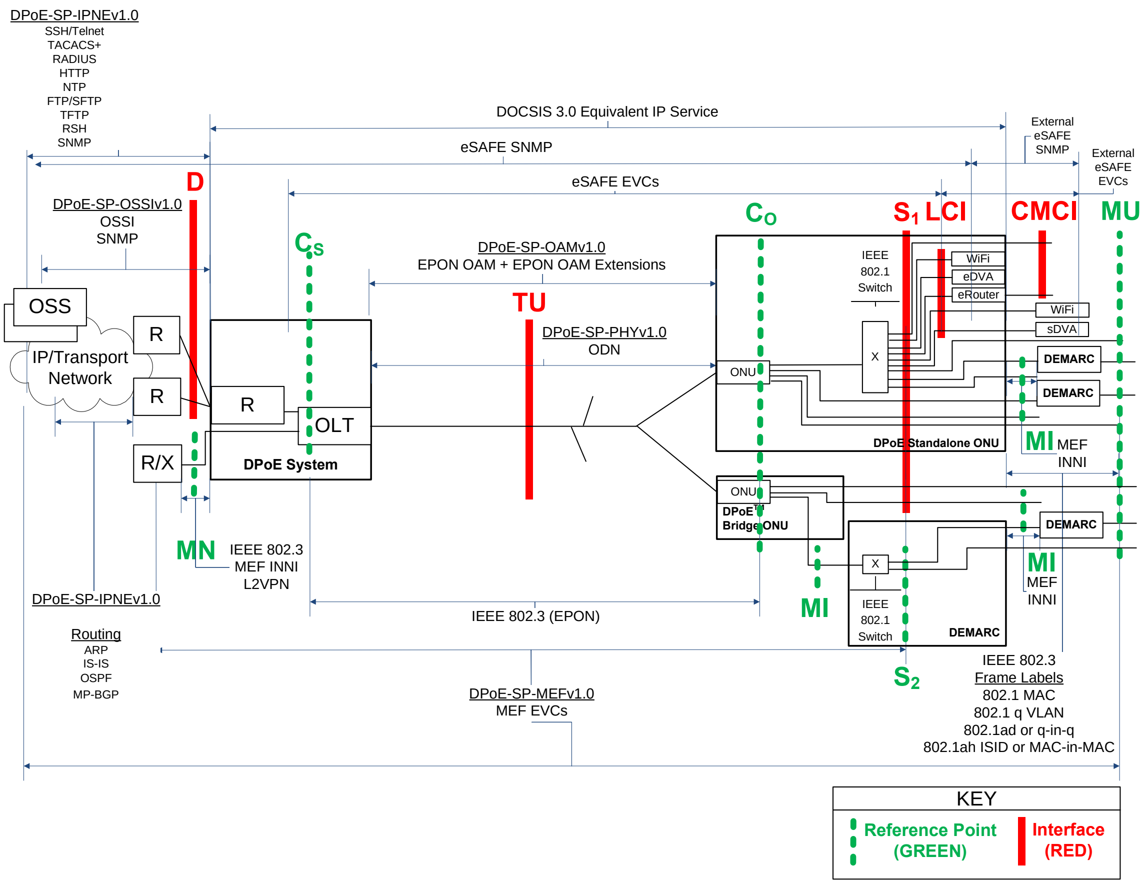

Figure 2 - DPoE Interfaces and Reference Points

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.

Interface or Reference Point		Interface or Reference Point Description
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 [DPoE-SP-MULPIv1.0].
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 References

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

[802.1] Refers to entire suite of IEEE 802.1 standards unless otherwise specified.

[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 802.1d™-2004, IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges

[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 802.3-2008, Carrier Sense Multiple Access with Collision Detection (CSMA/CD) access method and Physical Layer specifications, January 2008.

[802.3ah] IEEE 802.3ah™-2004: Amendment to IEEE 802.3™-2005: Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks, now part of [802.3].

[802.3av] IEEE 802.3AV-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.

[CMCIv3.0] Data-Over-Cable Service Interface Specifications, Cable Modem to Customer Premise Equipment Interface Specification, CM-SP-CMCIv3.0, 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-MULPIv1.0] DPoE-SP-MULPIv1.0, DOCSIS Provisioning of EPON, MAC and Upper Layer Protocols Requirements, Cable Television Laboratories, Inc.

[eDOCSIS] Data-Over-Cable Service Interface Specifications, eDOCSIS Specification, CM-SP-eDOCSIS, Cable Television Laboratories, Inc.

2.2 Informative References

This specification uses the following informative references.

[802.1ad] IEEE Std. 802.1ad-2005™, IEEE Standard for Local and Metropolitan Area Networks – Virtual Bridged Local Area Networks Amendment 4: Provider Bridges, May 2006.

[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].

[802.3ag] IEEE Std. 802.3ag™-2007, IEEE Standard for Local and Metropolitan Area Networks-Virtual Bridged Local Area Networks-Amendment 5: Connectivity Fault Management, January 2007.

[802.3as] IEEE Std. 802.3as-™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.

[DOCSIS] Refers to DOCSIS 3.0 unless otherwise specified.

[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.

[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.

[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.

[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.

[MULPIv3.0]	Data-Over-Cable Service Interface Specifications, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0, Cable Television Laboratories, Inc.
[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 2011]	IETF RFC 2011, SNMPv2 Management Information Base for the Internet Protocol using SMIv2, K. McCloghrie, November 1996.
[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), R. Presuhn, Ed., 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.
[RFC 1918]	IETF RFC 1918, Address Allocation for Private Internets, February 1996.
[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).

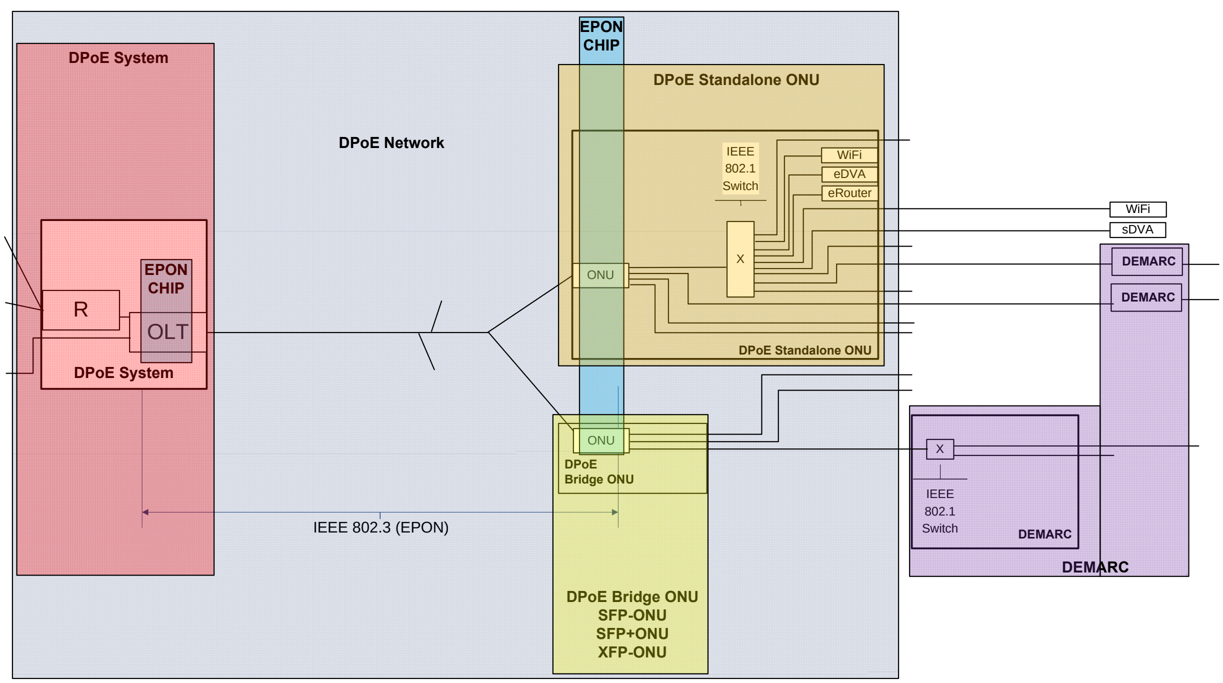

Figure 3 - DPoE Elements

3.2 Other Terms

1G-EPON	EPON as defined in [802.3ah]
10G-EPON	EPON as defined in [802.3ah] and amended in [802.3av]
Cable Modem CPE Interface	CMCI as defined in [MULPIv3.0]
Customer Premise Equipment (CPE)	Customer Premise Equipment as defined in [DOCSIS]
Multi-Layer Switching (MLS)	A switch that can switch based on Layer 2, Layer 3, Layer 4, etc.
Ethernet Passive Optical Network (EPON)	Refers to both 1G-EPON and 10G-EPON collectively
EPON Operations and Maintenance Messaging (OAM)	EPON OAM messaging as defined in [802.3ah] and [DPoE-SP-OAMv1.0]; Ethernet OAM is not the same as EPON OAM; Ethernet OAM is [802.1ag]
Logical CPE Interface	LCI as defined in [eDOCSIS]
Network Interface Device (NID)	A DEMARC device in DPoE specifications

4 ABBREVIATIONS AND ACRONYMS

This specification uses the following abbreviations:

CMCI	Cable Modem CPE Interface
CMIM	Cable Modem Interface Mask
CPE	Customer Premise Equipment
CoS	Class of Service
CTBH	Cell Tower Backhaul
CWDM	Coarse Wavelength Division Multiplexing
DIA	Dedicated Internet Access
DPoE	DOCSIS Provisioning of EPON
DR	Default Router
DSx	Digital Signal (DS1 or DS3)
DWDM	Dense Wavelength Division Multiplexing
eCM	embedded Cable Modem
eDVA	embedded Digital Voice Adapter
ENNI	External Network to Network Interface
EPL	Ethernet Private Line
EPON	Ethernet Passive Optical Network
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
Gbps	Gigabits per second (as used in the industry)
GBd	Gigabaud
IP(HSD)	High Speed Data (Broadband Internet Access using DOCSIS)
INNI	Internal Network to Network Interface
IP	Internet Protocol
I-SID	[802.1ah] I-Component Service IDentifier
LCI	Logical CPE Interface
LLID	Logical Link IDentifier
MEF	Metro Ethernet Forum
MEN	Metro Ethernet Network
MI	MEF INNI Interface at a customer premise
MLS	Multi-Layer Switching
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
OAM	EPON 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
PB	Provider Bridging [802.1ad]
PBB	Provider Backbone Bridging [802.1ah]
P2P	Point-to-Point
P2PE	Point-to-Point Emulation
P2MP	Point-to-Multi-Point
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
ROADM	Reconfigurable Optical Add-Drop Multiplexer
RS	Reconciliation Sublayer
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
WDM	Wavelength Division Multiplexing
WSC	Wireless Switching Center
X	IEEE Ethernet Switch (Generic)
XFP	X Form-factor Pluggable

5 DPOE SERVICE REQUIREMENTS

The purpose of this section is to document service requirements in sufficient detail to justify specific technical requirements for DPoE specifications. This document and this section do not attempt to identify the size of the markets, explain the business interests, or promote one or more services above any others.

Neither this document nor this section explains the variations of product feature sets required for a specific customer. Again, the focus is on technical requirements. For example, one operator might desire to have a single box ONU solution that fulfills all Ethernet business service customers. Another operator might prefer one ONU product for Ethernet customers below 1Gbps and another for customers above 1Gbps. Other variations might include combinations with other service sets (like IP services, voice, video, etc.), variations in data rate capabilities, variations in packaging and environmental requirements (wall mount, indoor versus outdoor, etc.). None of these variations is explored here. The services requirements described here are those that have an impact on the Operations, Administration, Maintenance, and Provisioning, and data collection for a multi-vendor EPON environment with DOCSIS, IP, and Ethernet controls.

Most of these services require core functionality that is shared across multiple services. Some require additional service-specific functionality.

5.1 Business Services Overview

The majority of business services can be delivered using Ethernet. In the access network, Ethernet can be used to deliver native Ethernet services (as defined by the Metro Ethernet Forum or MEF), or IP services transported over Ethernet. IP can, in turn, be used to deliver private IP (IP-VPN) services or public IP (Internet) services. IP can be used to deliver some storage network services, VoIP, or pseudowire TDM emulation. Pseudowire circuit emulation can be used to provide private line (T1 or T3 for example), IP-PBX, cell tower backhaul, etc.

Ethernet is the least common denominator of all modern telecommunications services. Whether in-house carrier services, wholesale, business services, or even residential, Ethernet has become the foundation for telecommunications services.

EPON provides a native Ethernet service and is an ideal access and transport technology. Accommodations for IP services provide network and management efficiency. The centralized control of the OLT provides tight scheduler control of both downstream and upstream traffic. EPON is capable of delivering any Ethernet or Ethernet-based service that conforms to IEEE 802 standards.

5.2 Service Requirements

Business services that do now, or may in the future, require EPON-based access/transport include:

- Private Data Services
    - Ethernet Services
- IP (Public IP/Internet)
- Voice Services
- Vertical Markets
    - IP-PBX
    - T1 (Pseudowire)

    - Cell Tower Backhaul
        - Ethernet or
        - T1 and T3 (Pseudowire)
    - Video Distribution and Transport (internal or external)
    - Storage network services
- In-House Use
    - Data center applications
    - Headend / hub applications
    - Transport between headends and hubs

These services are often complex because it is possible to combine and package these services in a variety of permutations and combinations. For example, a voice service may, to the access network, be simply a VLAN for an operator offering Ethernet transport based IP-PBX services. The same could be true for CTBH. Services can also be combined. For example, one customer might have private Ethernet transport, Internet access, private IP, and multiple voice services (IP trunking, pseudowire, and PSTN services) or any combination in one or more sites.

5.3 Single Tenant Businesses

Single tenant businesses can be served with a single dedicated ONU. Each business customer may have one or more services. A 'service' in DPoE specifications is the network service provided on a single UNI port, such as CMCI for IP(HSD) or MU for MEF, or an LCI for an embedded Service/Application Functional Entity (eSAFE). The analogy for an LLID in DOCSIS specifications is the DOCSIS SF. The DPoE System MUST dedicate one LLID for each service. Similarly, the DPoE ONU MUST use the same dedicated LLID for that service. The LLID is the baseline technology for the guaranteed delivery of each individual service delivered over a DPoE Network.

Since each separate service has separate requirements, service delivery must be guaranteed for each individual service. Probabilistic techniques that are not absolutely controlled or controllable are not sufficient to meet this requirement. Therefore, any QoS technologies that use packet or frame marking and queuing only (without a fixed-time algorithm scheduler) are not acceptable. In particular, the LLID scheduler for EPON is a TDMA algorithm, which guarantees service delivery for each properly configured LLID. Packet and frame marking technologies are still both useful, and required, for traffic management within an individual service; however, this requirement is distinguished from the need to guarantee individual services.

5.4 Multiple Tenant Businesses

Multiple tenant buildings will differ from single tenant business buildings or mixed-use residential and business buildings. Although it is possible, it is not highly practical to operate a multiple tenant business service by using multiple ONUs and fiber splits within the facility.

The most cost-effective use of fiber would utilize a single strand into a building with a single ONU to service multiple customers. The product variations required for different sized facilities, with variations in service combinations, strongly suggest a modular solution. However, identifying the correct product is not the objective of this effort. Rather, the objective is to identify the baseline requirements for provisioning and operating such a solution. The most flexible solution for both single tenant and multiple tenants will, in some cases, be a modular approach. Such a modular approach could include a DPoE Standalone ONU and a separate DEMARC device or a DPoE Bridge ONU plugged into a DEMARC device.

5.5 Detailed Service Requirements Introduction

Each group of services below can be further broken down into more detailed services that are closer to the product descriptions of individual services that operators offer to their customers.

- Private Data Services
  - Ethernet Service
    - E-LINE
      - EPL
      - E-VPL
    - E-Tree
      - EP-Tree
      - EVP-Tree
    - E-LAN
      - EP-LAN
      - EVP-LAN
  - Private IP (IP-VPN)
- IP (Public IP/Internet)
  - IP(HSD) Fixed IP/Internet Service (CMTS IP Routing)
    - Static Routing with 1 fixed IP address/MAC
    - Multiple fixed IP address/MAC
  - Internet Service
    - RIP based address learning
    - BGP based address learning
- Voice Services
  - May include IP-PBX services (described below) but categorized as voice
  - Single line telephony
  - Multi-line telephony
- Vertical Markets
  - Wholesale Ethernet
  - T1 or T3 Pseudowire
  - IP-PBX (Can be implemented as an Ethernet, Private IP, Public IP, or "Voice" service)
  - Cell Tower Backhaul
    - Ethernet
    - T1 or T3 Pseudowire
  - Video Distribution and Transport (internal or external)

Multicast services within Metro Ethernet services, or for the IP services, are not explicitly considered in this version of the specifications. Although IP multicast services can be operated on top of any EPON service, just as they can be operated on top of any Ethernet technology, EPON has unique capabilities to introduce greater broadcast and multicast efficiencies because it is a P2MP technology that natively supports broadcast. IP multicast services could be manually provisioned and operated by operators on a DPoE Network. However, this version of the DPoE specifications does not consider the operations, administration, maintenance, or provisioning of IP multicast.

5.6 Private Data Services

5.6.1 Ethernet Service

EPON is a native Ethernet transport for Ethernet service or Metro Ethernet services. Services provided on top of Ethernet can be constructed in a variety of ways. For example an IP service can be constructed by building an Ethernet service from the ONU UNI (the DPoE S interface) to a MEF UNI terminated at an IP router (physical or sub-interface).

Additional variations might include technology variations for operator needs in servicing these requirements. For example, if an operator has two customers in two separate buildings, their needs might be different from two similar customers co-located in a single building. This has a direct correlation to EPON requirements. A single tenant solution with a single ONU might include several services with several ports and thus require a small number of LLIDs. For a multi-tenant solution, it might be necessary to use a DPoE ONU that has more ports to accommodate a larger number of customers. In that case, the number of required LLIDs for the multitenant ONU is expected to be significantly larger.

5.6.2 Private IP

IP service can be constructed with EPON by building an Ethernet service from the ONU UNI (the DPoE S interface) to a MEF UNI terminated at an IP router (physical or sub-interface).

The Ethernet Service could be constructed to another customer or customer site on EPON or any other access technology. That Ethernet service could also be constructed to reach service platforms within the operator's network. Examples include a Class 5 server (switch) for IP-PBX or VoIP-based voice service, an IP router for BGP peering, IGP (RIP or static) routing, an IP router for IP-VPNs, Ethernet switches or virtual Ethernet switches (VPLS, H-VPLS, or Layer 2 bridge groups) for Ethernet services, etc. This is the architecture already widely adopted by operators for converged service transport.

5.7 Internet Service

Internet service can be offered by two different methods in the DPoE specifications. DPoE specifications support both DOCSIS IP(HSD) equivalent service and MEF-based transport to emulate what operators have sometimes called Dedicated Internet Access (DIA).

5.7.1 Static Address Service

5.7.1.1 IP(HSD) Method

Internet service can be offered by two different methods in DPoE specifications. Static address-based services can be implemented in exactly the same manner as with DOCSIS 3.0 services, utilizing either a bridged service to the CMCI interface or the LCI interface with an eRouter as illustrated in Figure 1 and Figure 2.

Since the DPoE architecture is based on Metro Ethernet services as defined by MEF, it is possible to extend the EVC beyond a DPoE ONU by creating another MEF trail from a DPoE LCI interface to a DEMARC device (directly attached), or a remote DEMARC device (reachable across an Ethernet network). The MU interface on a DEMARC device or MEF UNI interface anywhere in a remote Ethernet network terminates the MEF EVC, providing the DOCSIS CMCI interface at such a remote location. Although this method is supported, this version of the DPoE specifications does not specify requirements beyond the DPoE Network. Configuration of the DEMARC device, and any MEN between the DPoE ONU and the MEF UNI, is out of the scope of this specification.

5.7.1.2 Ethernet Method

The operator can also provide a static IP address service using [DPoE-SP-MEFv1.0] to provide Ethernet service between the customer equipment and another MEF UNI interface connected to the operator's IP router. While the operator's router for this service may be connected to a MEF UNI interface, the interface on the operator's router does not have to be a DPoE interface.

This service could be implemented with a wide variation of Metro Ethernet service types. For instance, at the MU connected to the operator's default router, the operator could construct a series of EVPLs where the far-end of each EVPL connects to a different customer or customer location, some of which may be on a DPoE Network. Alternatively, an E-LAN could be used to provide a functionally equivalent IP(HSD) service.

5.7.2 Dynamic Address Service

Dynamic customer address learning via an IGP (such as RIP) can be supported using the same method, where the DPoE S interface is used to provide Metro Ethernet service to an operator router (the DEMARC device) at the customer premise. It can also be supported by using a DPoE Bridge ONU to provide Metro Ethernet service to an operator router that acts as the DEMARC device.

5.7.3 Transit Service

Internet transit service using BGP can be provided using the same Metro Ethernet services specified in [DPoE-SP-MEFv1.0].

5.7.4 Peering Function

Operators could use DPoE Networks for peering with (non-paying) Internet peers in exactly the same manner as for transit service.

5.8 Voice Services

Voice services are not directly supported in this version of the DPoE specifications. VoIP applications other than PacketCable 2.0 can be supported with VoIP or DS1 pseudowire 'over the top' using Ethernet or IP over Ethernet services of DPoE.

5.8.1 Single Line Telephony

PacketCable voice service is not directly supported in this version of the DPoE specifications.

5.8.2 Multiple Line Telephony

PacketCable voice service is not directly supported in this version of DPoE specifications.

5.9 Vertical Markets

A vertical market is a market for telecommunications services that requires specific services or functions in the network services or management of the network services. Vertical markets often require additional functionality such as specialized network management services, on-demand or custom provisioning capability, and non-standard network interface support from customer premise equipment.

DPoE specifications do not contain any requirements that were designed specifically for a vertical market. However, the architecture for DPoE Networks was designed to allow 'over the top' management of services either within a standalone DPoE ONU or beyond the DPoE S interface or reference point on the far side of a DEMARC device as illustrated in Figure 1 and Figure 2.

Support for over the top control, signaling, and management is provided by EVCs across the C_O reference point for DPoE Bridge ONU implementations. Such a management connection is no different than other EVCs.

For a DPoE Standalone ONU, there are three possible models for signaling and management. The first is an embedded model. In the embedded model, an EVC extends from the DPoE System to an eSAFE within an eDOCSIS device (a DPoE ONU) across the LCI interface. This EVC is terminated on the DPoE System. In the second model, the operator would construct an over-the-top EVC through the S_1 interface where there is a single EVC on a single S_1 interface. In the third model, the operator would construct an over-the-top EVC within an ELINE or EVPL through an S_1 interface. Such a model could, for example, be used to provide a management connection (such as a VLAN) to a DEMARC device over the same S_1 interface carrying bearer traffic as illustrated in Figure 2. In both over-the-top models the constructed EVC could terminate on some other device, such as an MU on another DPoE system, which connects to some operator's service management network.

One or more EVCs can be constructed to allow either combined or separate logical circuits for signaling, management, and bearer traffic. Those EVCs can be provisioned across the MEN to another customer on EPON or any other access technology. The EVC could also be provisioned back to service platforms within the operator's network. Examples include a Class 5 server (switch) for IP-PBX or VoIP-based voice service, an IP router for BGP peering, IGP (RIP or static) routing, an IP router for IP-VPNs, Ethernet switches, or virtual Ethernet switches (VPLS, H-VPLS, or Layer 2 bridge groups) for Ethernet services, etc. This is the architecture already widely adopted by operators for converged service transport.

5.9.1 Wholesale Ethernet

The service requirements for wholesale Ethernet in the access network are defined by MEF. Wholesale Ethernet differs from direct business services only in back office operations and in inter-carrier Network to Network Interface (NNI), which MEF calls the External NNI (ENNI). The current [MEF 26] specifications are based on manually mapping Ethernet [802.1ad] tags between MEF Metro Ethernet Networks (MEN). Like DPoE, [MEF 26] is organized into phases that will deliver progressively more complex services. [MEF 26] is the Phase I ENNI for MEF. Carrier interconnection is outside of the scope of the access network, and therefore, not defined by DPoE.

EPON is well-suited to wholesale Ethernet. It can provide Metro Ethernet services with granular control as accurate as single 1kbps in each direction on each EVC. As an [802.3] Ethernet protocol, 1G-EPON and 10G-EPON cannot transport frames larger than those permitted by [802.3] of 2000 bytes[1]. Because 1G-EPON was developed prior to the Amendment of [802.3] based on [802.3as] (for 2000 bytes), most 1G-EPON implementations are limited to 1600 bytes. The frame size requirements for DPoE Networks are specified in [DPoE-SP-MULPIv1.0].

[1] IEEE 802.3 support for 1518 byte payload remains the same in both 1G-EPON and 10G-EPON. The additional 82 bytes (1600 byte total) in 1G-EPON or additional 482 bytes (2000 byte total) in 10G-EPON are for frame overhead. Frame overhead is used for IEEE 802.1ad, IEEE 802.1q, IEEE 802.1ah and other Ethernet tag and frame overhead.

5.9.2 T1 or T3 Pseudowire

DSx emulation over IP (over Ethernet) can be utilized to provide T1 or T3 access services. Latency and variation in latency over time (jitter) in EPON are well-bounded by the maximum delay in the OLT scheduler (per PON). Although there are no EPON-specific standards for DSx emulation over EPON, any DSx pseudowire solutions for Ethernet or IP can be transported over EPON.

T1 can be used for variety of applications ranging from customer service-agnostic private line to PBX extensions, carrier wholesale, cell tower backhaul, or other specific services.

Although jitter in DSx emulation is minimal, pseudowire solutions are not accurate enough to meet the requirements of ITU-T G.824, which specifies the DSx hierarchy for timing signals. In particular, the lack of an absolute synchronous signal leaves wander as an outstanding problem for pseudowire solutions. This challenge is not specific to EPON.

EPON is capable of native T1 transport without IP. Although some products support DSx over EPON, neither EPON nor DPoE Networks have specific DSx standards at this time. DPoE specifications do not preclude such support, which may be implemented provided that the operation of such services does not affect DPoE operation, including the management and provisioning of both ONU-specific and total PON bandwidth and scheduling (QoS or CoS) specifications and requirements.

5.9.3 IP-PBX

IP-PBX services can be supported just like any other IP service. Although there are no specific provisions for IP-PBX, there are a variety of methods to support IP-PBX over EPON or DPoE. Both require the manual provisioning of EPON bandwidth. IP-PBX could be implemented by:

- IP through LCI interface over EPON transport through the D interface to a remote IP-PBX
  - Using the DOCSIS 1.1 IP service solution (bridged to DPoE System router).
  - Using the DOCSIS 3.0 IP service solution (bridged or routed with eRouter on DPoE Standalone ONU to DPoE System router).
- MEF-based Ethernet transport solution
  - Using the MEF solution to transport an EVC from a call platform across the DPoE Network and through an MI interface to a DEMARC device with an IP-PBX or with an IP-PBX attached.
  - Using the MEF solution to transport an EVC from a call platform across the DPoE Network and through an MU interface to an IP-PBX attached.

The latter is possibly the most flexible option. An EVC can be constructed from the S interface to a Class 5 server (switch) operating the IP PBX service. Such services have already been successfully deployed by operators using EPON and can be supported by [DPoE-SP-MEFv1.0].

5.9.4 Cell Tower Backhaul

Cell tower backhaul (CTBH) requires either or both DSx and Ethernet services. DSx CTBH can be supported using DSx pseudowire, non-standard DSx over Ethernet, or DSx over EPON solutions as described above.

An EVC can be constructed from the MU (or via an S interface to a DEMARC providing an external MEF UNI) interface to a Mobile Switching Center (MSC) or Wireless Switching Center (WSC). For Ethernet, the service can be delivered as an EVPL to meet MSC or WSC Ethernet aggregation requirements. For DSx emulation, the EVPL can terminate at a pseudowire service aggregation device. Such services have already been successfully deployed by operators using EPON and can be supported by [DPoE-SP-MEFv1.0].

5.9.5 Video Distribution and Transport

Video can be transported over EPON, within a facility or between facilities, using video over IP or video over Ethernet. Native video transport over Ethernet is compatible with current Gigabit Ethernet video technologies. Video distribution using EPON within a facility does not require DPoE specifications because it will typically not be a dynamic access configuration environment, but rather a static provisioned implementation.

Video transport over EPON for third parties can be supported by using IP(HSD) forwarding or MEF forwarding as specified in [DPoE-SP-MULPIv1.0] and [DPoE-SP-MEFv1.0].

5.9.6 Storage Network Services

DPoE specifications do not contain any specific requirements or recommendations for storage network services. As with all EPON services, EPON can only support IEEE standard frame sizes. As an IEEE 802 Ethernet protocol, 1G-EPON and 10G-EPON cannot transport frames larger than those permitted by [802.3] of 2000 bytes[2]. Because 1G-EPON was developed prior to the Amendment of [802.3] based on [802.3as] (for 2000 bytes), most 1G-EPON implementations are limited to 1600 bytes. The frame size requirements for DPoE Networks are specified in [DPoE-SP-MULPIv1.0].

Jumbo frames are often requested by customers for Storage Area Networks (SANs). EPON cannot support jumbo frames. EPON maximum frame sizes are based on [802.3] specifications. Although it is possible to design an EPON system that uses larger frame sizes, such a system would go beyond the [802.3] specifications and potentially create additional complexity, such as the increase in jitter that would come with greater variation in frame size.

Some storage services such as iSCSI do not require (although they optionally support) jumbo frames. For these services, EPON can be used for storage network transport by delivering Ethernet with [DPoE-SP-MEFv1.0].

[2] IEEE 802.3 support for 1518 byte payload remains the same in both 1G-EPON and 10G-EPON. The additional 82 bytes (1600 byte total) in 1G-EPON or additional 482 bytes (2000 byte total) in 10G-EPON are for frame overhead. Frame overhead is used for IEEE 802.1ad, IEEE 802.1q, IEEE 802.1ah and other Ethernet tag and frame overhead.

6 ARCHITECTURAL FOUNDATION

6.1 Forwarding

The foundation for all DPoE services is an Ethernet service model as described by MEF. Both the Metro Ethernet services and IP services are based on a common Metro Ethernet service architecture.

6.1.1 Ethernet Virtual Connections

The basic unit of services in this version of DPoE specifications is an Ethernet Connection as defined by MEF. MEF defines an Ethernet Virtual connection (EVC) between user-defined end points called User to Network Interfaces or UNI. As in MEF, the EVC is the least common denominator of service for DPoE.

In [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 by DPoE, for example, the S-VLAN and C-VLAN are used as "outer" and "inner" tags (like q-in-q) without respect to their typical meaning in [802.1ad].

A DPoE Network can be used to provide a MEF TRAN-trail in the First Mile or Last Mile or both. In this version of DPoE, First and/or Last Mile provisioning are independent of each other.

IP services are implemented over the top of Metro Ethernet services to allow for:

- More than one CMCI interface on a DPoE ONU
- Uniform forwarding and multiplexing model
- Uniform QoS rules in the access network
- Extensible architecture beyond the access platform into the operator MEN
- Extensible architecture beyond the access platform into a customer premise MEN
- Interoperability with other MEF systems
- Allow for remote DOCSIS devices by extending the [eDOCSIS] logical bridged Ethernet [802.1] model (LCI interface) beyond the "CM" (or in this case the DPoE ONU) to external [802.1] bridges across a common Metro Ethernet service model (which is also based on [802.1] forwarding). Each LCI interface MUST support exactly one eSAFE (either embedded in the ONU or external).

For some services, the EVC may terminate at a service solution within the hub. For other services, the EVC may traverse a Layer 2 transport (what MEF calls a Metro Ethernet Network or MEN) to a remote UNI within the operator's network or at another customer premise. On top of each EVC, additional services such as IP can be implemented. In other words, the fundamental building block of services delivered with DOCSIS Provisioning of EPON will be an EVC from (either a DEMARC device to) a DPoE ONU to a DPoE System. From the DPoE System, the DPoE EVC trail can be "spliced" into Ethernet transports (real or virtual) or "terminated" at an IP router interface for IP transport and services. When an IP routing function is operated on the DPoE System, the IP service can be provisioned by DPoE. When an IP routing function is not operated on the DPoE System, the IP service cannot be provisioned by DPoE.

6.1.2 Metro Ethernet Services

Metro Ethernet services require a variety of architectural needs. These include single box MEF solution, support for external NIDs (which we call DEMARC) devices, and many types of INNI interfaces at the MN. Each service that MEF describes needs to be supported over each of these different models. MEF itself supports [802.1d], [802.1q], q-in-q, [802.1ad] (C tag only), and [802.1ad] (Provider Bridged) frames. MEF also offers a variety of QoS parameters.

DPoE Networks are capable of supporting all Metro Ethernet services. Full DOCSIS service automation is limited by the work required to develop the service models and automate DOCSIS file-based provisioning and operations. The specifications for Metro Ethernet services are described in [DPoE-SP-MEFv1.0].

6.1.3 Internet Services

DPoE Networks support MEF UNI based EVCs, but extends the use of EVCs not only for UNI interfaces, but also for soft UNI interfaces. A soft UNI interface is an interface that acts like a MEF UNI interface, but which is not a UNI interface. DOCSIS IP services, for example, are implemented using an EVC from a DPoE ONU Ethernet port to an IP Router function within the DPoE System. The DOCSIS IP service (sometimes called High Speed Data or IP(HSD) is a broadband Internet access service. It is implemented using DPoE specifications in one of two ways. DOCSIS 1.1 equivalent service is a "bridged" service from the customer premise through a DPoE ONU LCI interface to a default router (DR) operating on a DPoE System. The "bridge" (as it is called in DOCSIS specifications) is implemented in DPoE Networks as an EVC from the DPoE ONU LCI interface to the DPoE System. The second implementation is [eRouter], which operates an IP router within the DPoE ONU. In this implementation, the EVC goes from the router in the DPoE System to an LCI interface on an eRouter (which is a type of eSAFE) in a DPoE ONU that is operating as an eDOCSIS CM.

Each DOCSIS IP serving group (IP-SG as identified in [DPoE-SP-MULPIv1.0]) can be associated with a single [802.1ad] S-VLAN. Each customer is mapped into a single [802.1ad] C-VLAN ID (VLAN).

DPoE Networks are capable of supporting all DOCSIS IP-based services. One fundamental distinction between DOCSIS and DPoE Networks is that support for Ethernet services in DPoE Networks is native, and does not need to be tunneled or emulated. IP-based encapsulation for Ethernet service transport is, therefore, not required. The IP(HSD) and VoIP IP services can be implemented over EVCs as described above. The specifications for IP services are described in [DPoE-SP-MULPIv1.0].

6.2 Service Integration

In DOCSIS the service provisioning and services are tied closely together. This remains the case with DPoE specifications. DPoE specifications use the EPON LLID as the primary mechanism for mapping Metro Ethernet service instances to the MAC and MAC sub-layers in order to provide consistent Quality of Service based on the TDM capabilities of EPON. An EPON LLID is logically equivalent to an SF in DOCSIS. IP(HSD) services in DPoE Networks always operate over Metro Ethernet service, and therefore IP(HSD) also uses this common model.. This is accomplished by mapping each Metro Ethernet service EVC to one LLID. This provides a simple, but scalable, architecture limited by the available number of LLIDs on a DPoE ONU or on each EPON.

6.3 Multipoint Architecture

There is one distinct way that EPON is very different from DOCSIS and Ethernet. That difference is derived from the multi-access architecture of EPON. EPON is fundamentally a multi-access optical fiber network. Although it can be used to provide broadcast and unicast services just like DOCSIS, the multi-access function in EPON does not operate as it does in DOCSIS. Broadcast services in EPON are natively supported at the physical layer due to the properties of the underlying physical fiber plant (passive forward split optical distribution network). EPON supports both forward broadcast and unicast when using the MAC. Unlike DOCSIS, return transmissions are non-broadcast. This difference is because the fiber architecture for a PON passively optically combines the return path signals, over time, in the feeder portion of the ODN. The EPON TDMA implementation features a master scheduler in the OLT to time division multiplex upstream burst laser transmissions. The portion of fiber in the ODN from the splitter to the customer premise is a dedicated access medium that is a non-broadcast transmission path in the reverse direction. However, in DOCSIS, the CMs must deal with the analog artifacts of TDMA on a shared electrical medium. EPON ONUs transmit on a dedicated fiber drop that appears to be a dedicated return channel. While it is dedicated from the split to the ONU, it is shared, beginning at the splitter. EPON's Multi-Point Control Protocol (MPCP) handles the management of the shared return path. The OLT operates as a master scheduler, and the ONUs are slaves to the OLT. This is possible because the fundamental technology of current PONs is the burst-mode laser. The laser at the ONU can be turned on and off for short bursts of transmissions.

EPON's MPCP, in combination with the service control capabilities of the Logical Link Identifier (LLID), offers the capability to emulate everything that is done in DOCSIS today, add on native Metro Ethernet services, and provide managed service for broadcast (multicast) because EPON is not an asynchronous [802.3] transport. MPCP is a true multi-layer protocol specifically designed to support a single control plane for unicast Ethernet or IP services, broadcast Ethernet services, and multicast IP services. Multicast is supported by physically broadcasting frames on the PON that are received by ONUs.

6.4 DOCSIS Emulation

A DPoE Network, to outside systems (CPE and the operator OSS), acts like a DOCSIS system with a CMTS and attached CMs. IP interfaces on the DPoE System act like those on a DOCSIS CMTS. The DPoE ONU provides the same service capabilities as a CM.

Most ONUs do not have an IP software stack and lack the resources to operate as a CM. The lack of an IP stack is both an additional security advantage and an economic advantage of EPON technology. Operators specifically want to avoid bifurcating the EPON ONU market and specifically want to take advantage of the benefits of scale in manufacturing and support of existing EPON ONU products and technologies.

Since all traffic going to a DPoE ONU from the OSS passes through the DPoE System, it is possible for the DPoE System to perform the functions of a CM. DPoE Systems operate a virtual Cable Modem (vCM) for each and every DPoE ONU. The vCM is implemented in the DPoE System. It handles all of the OAMP functions for DOCSIS as described in [DPoE-SP-MULPIv1.0] and [DPoE-SP-OSSIv1.0]. The vCM can proxy requests, signaling, and messages to the DPoE ONU using EPON OAM messages, including at least those specified in [802.3ah] or [DPoE-SP-OAMv1.0].

The vCM model applies only to the DPoE ONU. eDOCSIS devices (if implemented) use an IP on the DPoE ONU for all other DOCSIS services beyond CMCI based services. This includes eDVA, eRouter, eDSG, WiFi, or any other eSAFE subsystem within a DPoE ONU. Standalone DOCSIS devices such as an sDVA use IP on each respective device. The vCM model is only used for DOCSIS CM management functions.

Customer-forwarded MEF or IP traffic does not pass through the vCM. The vCM is typically implemented in IETF [RFC 1918] non-globally routed (private) IP address space because the vCM address is used only internally within the operator's network.

6.5 eDOCSIS

[eDOCSIS] is a DOCSIS CM architecture for embedded subsystems within a CM. DPoE specifications allow for the optional support of [eDOCSIS]. [eDOCSIS] requires an eCM that operates as a DOCSIS CM. Because DPoE specifications require the CM management function within the vCM, only the classifier and forwarding functions in the DPoE ONUs need to be implemented in the eDOCSIS device (the DPoE ONUs). Any DPoE ONUs that include [eRouter], eDVA, or other eSAFEs is an [eDOCSIS] device. DPoE ONUs that are eDOCSIS devices implement eCM in the eDOCSIS device (the DPoE ONUs) as specified in [eDOCSIS]. The implementation of an eCM consists of the vCM operating within the DPoE System and the eDOCSIS functionality (ONU firmware and any eSAFE software) in a monolithic code image as specified in [eDOCSIS].

[eDOCSIS] implementation in [DOCSIS] consists of an embedded CM (called an eCM) within a CM, virtual interfaces between the eCM and eSAFEs, and each eSAFE as illustrated below (in Figure 4). An eSAFE is an embedded Service/Application Functional Entity. It is an embedded subsystem.

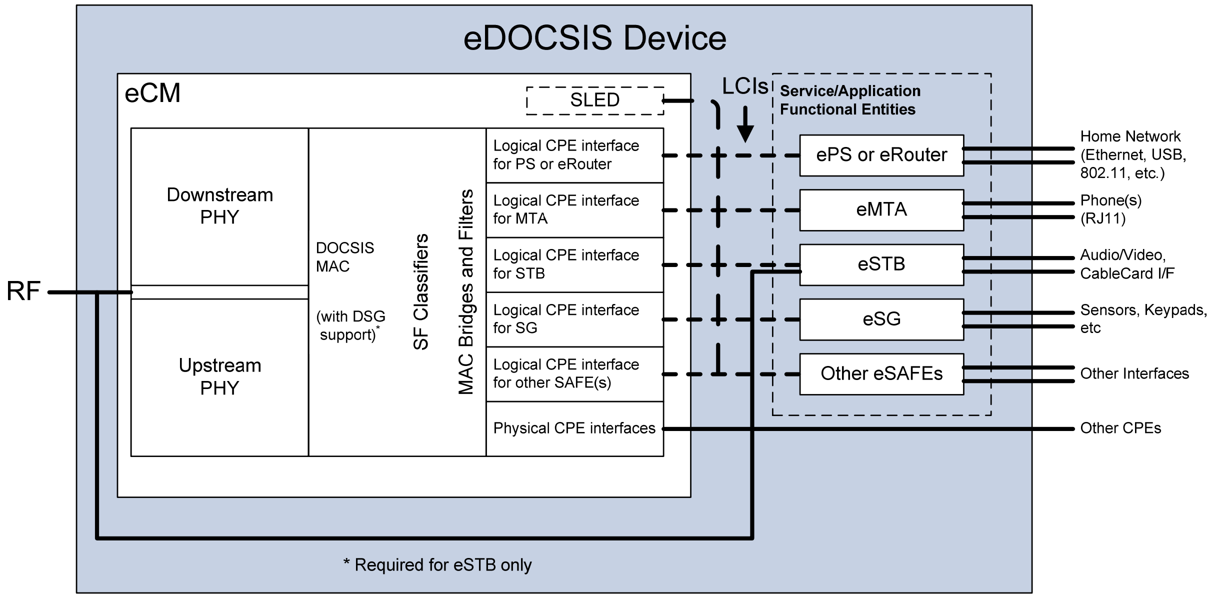

Figure 4 - eDOCSIS Reference Model from Figure 5-1 of [eDOCSIS]

The SLED interface (loopback function) is only required for eSAFEs that do not have physical interfaces on the eDOCSIS device. The eDOCSIS model further details the eCM as having PHY blocks (for downstream and upstream), a DOCSIS MAC, classifier, Ethernet [802.1] block, and Logical CPE Interfaces (LCIs).

Figure 5 shows a simplified and updated model of eDOCSIS. This model is simplified by showing the PHY as a single block, shown without SLED, and shown using an eDVA in place of the legacy eMTA.

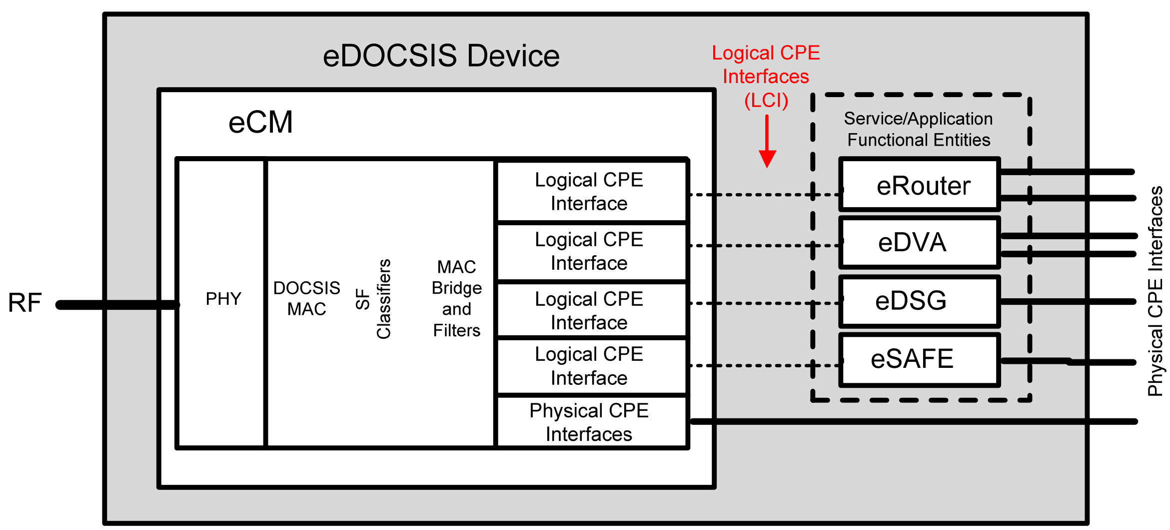

Figure 5 - Simplified and Updated Model of eDOCSIS

Further enhancing the model, Figure 6 simply adds a label for the SF classifiers (C_O) and replaces the RF interface with an Ethernet interface.

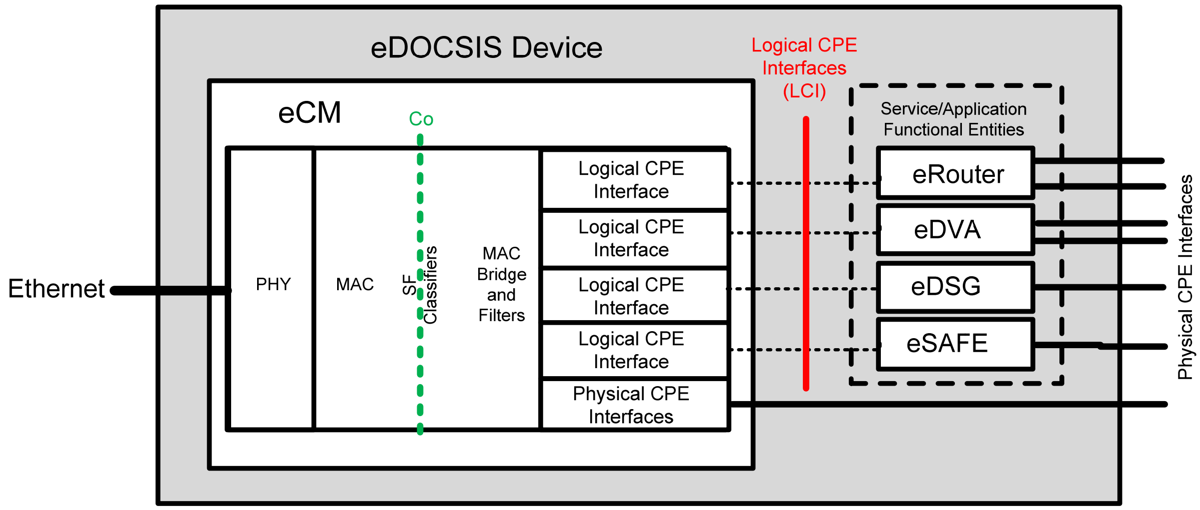

Figure 6 - eDOCSIS Model Based on Ethernet

In DPoE Networks, the MAC is the Ethernet MAC [802.3]. The forwarding function in the eDOCSIS device is based on [802.1]. The eCM functionality is replaced by the vCM functionality on the DPoE System (as described in Section 6.4 above). To keep consistency with both the DPoE architecture and the eDOCSIS architecture, the vCM can use EPON OAM and DPoE OAM Extensions to operate the eDOCSIS device in the same way as an eCM is operated.

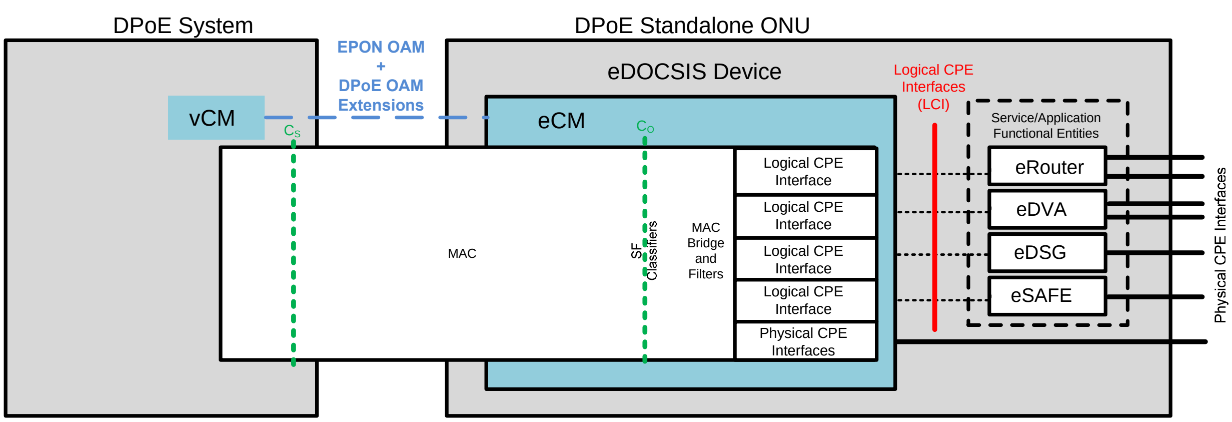

Figure 7 - DPoE eDOCSIS Adaptation

Although the EPON MAC contains an additional sublayer (MPCP), the Ethernet client MAC is preserved because the sublayer is transparent to the service layer. The MPCP control sublayer operates between the Ethernet MAC sublayer and the MAC client. The MAC in the DPoE eDOCSIS adaptation is the Ethernet MAC client. The "MAC Bridge and Filters" in eDOCSIS is really not a bridge in eDOCSIS (or in DPoE specifications). A bridge (by definition) is a two port system. A multi-port bridge is really a switch. Assuming an Ethernet forwarding model, it is an Ethernet switch. An Ethernet switch with "filters" is still an Ethernet switch. The filters are used for both upstream port (mask) matching as well as downstream port (mask) matching. The "mask" in DOCSIS is the Cable Modem Interface Mask (CMIM). The forwarding model for DPoE ONUs relies on the use of the CMIM with the use of classifiers that are extended (beyond the original DOCSIS classifiers) to include [802.1] headers for "filtering" and "classification." In DPoE specifications, the eDOCSIS "MAC Bridge," filtering, and forwarding are combined into a single forwarding model described in [DPoE-SP-MULPIv1.0].

The eDOCSIS model for DPoE specifications operates the eCM within the DPoE System as part of the vCM. What is left in the eDOCSIS devices is the Ethernet switch within the eCM. To comply with [eDOCSIS], the remaining eSAFEs and any DPoE ONU software are combined into a single software image as described below. Based on the eDOCSIS model, the eCM is split across the DPoE System and DPoE ONU.

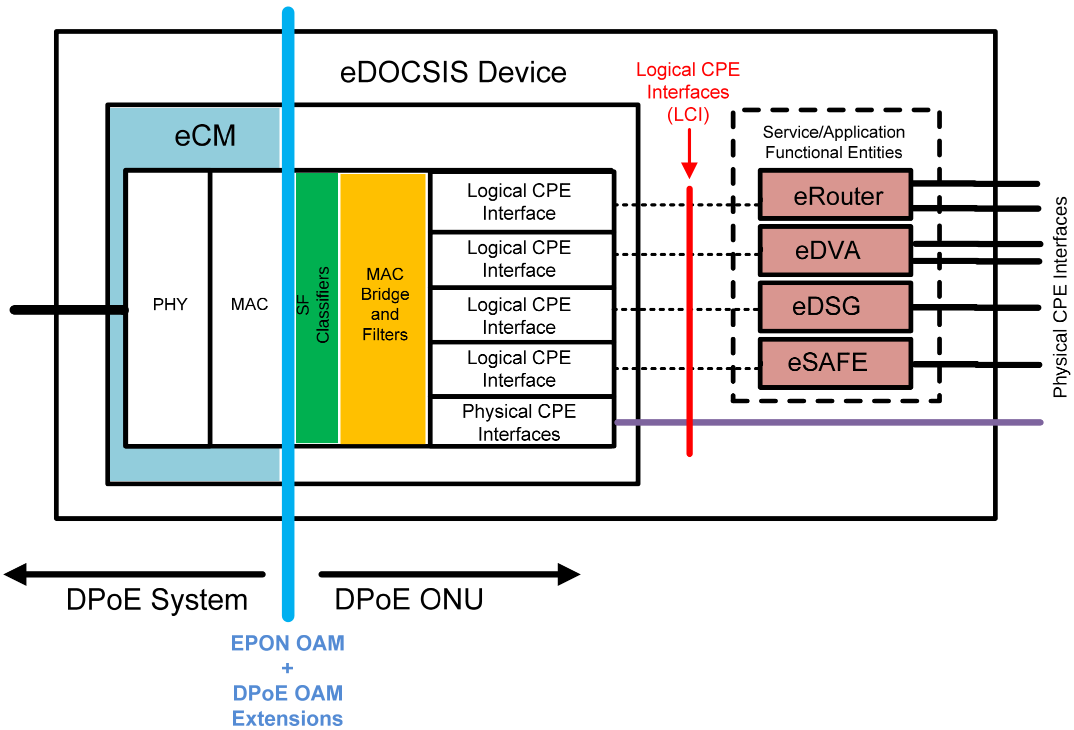

Figure 8 - DPoE Specifications Splits eDOCSIS

Each function within eDOCSIS maps to an equivalent function in the eDOCSIS architecture for DPoE as shown in Figure 9.

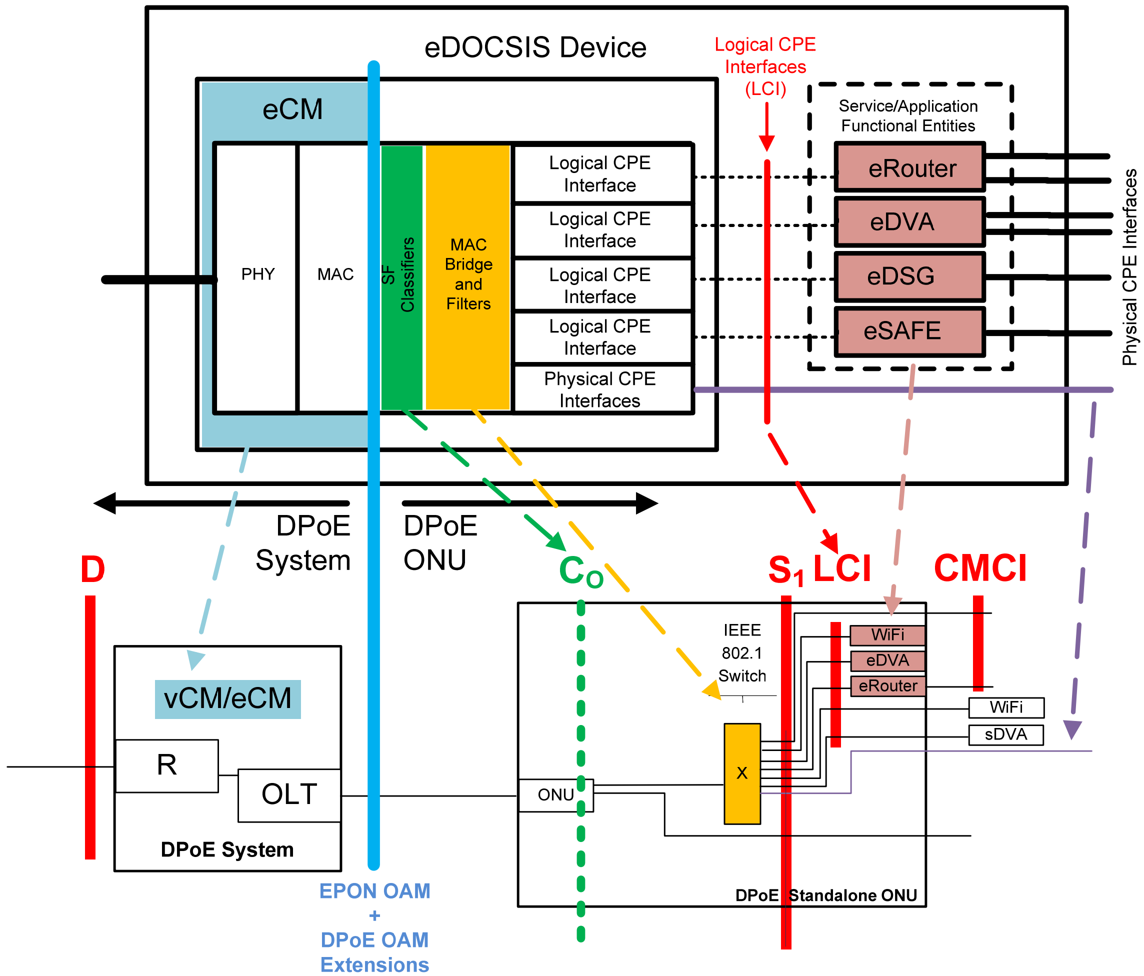

Figure 9 - eDOCSIS to DPoE eDOCSIS Mapping

The DPoE eDOCSIS architecture extends the eDOCSIS architecture to support standalone (external) devices that can act as part of the eDOCSIS device. This new capability allows operators to use external device management (such as WiFi, sDVA, DEMARC, or other devices) as part of the DOCSIS OSSI based configuration. In order to achieve this integration, the software for such systems could be packaged with the DPoE ONU's software to comply with the [eDOCSIS] specifications.

6.6 EPON OAM

As described above, the control path for communication between the DPoE System and DPoE ONUs are EPON OAM messages as defined in [802.3ah]. EPON OAM messages are datagrams (called Protocol Data Units or PDUs) sent between the OLT and ONU (as described in [802.3] and [DPoE-SP-OAMv1.0]. These messages are available only on the PON, and they are sent directly from an OLT to ONU or ONU to OLT. EPON OAM messages are not forwarded beyond the OLT or beyond the ONU.

EPON OAM is distinct from Ethernet OAM, such as [802.1ag] or MEF OAM as defined in [MEF 21].

DPoE Networks use [802.3ah] standard OAM messages and OAM Extensions (permitted by [802.3ah]) as specified in [DPoE-SP-OAMv1.0].

7 REQUIREMENTS

7.1 Introduction

The technical requirements for DPoE Networks are driven by three factors. The first sets of requirements are for back office compatibility with existing DOCSIS infrastructure and DOCSIS-based operations, administration, maintenance, and provisioning. The second sets of requirements are for specific services. The third set of requirements is built on support for existing [802.3ah] and [802.3av] specifications and products.

One very fundamental difference between DOCSIS and DPoE Networks is that the least common denominator of services in DOCSIS is IP, where the least common denominator of service in DPoE Networks is Ethernet. DOCSIS services assume an IP transport and must emulate Ethernet services. DPoE specifications assume an Ethernet service that can run Ethernet or IP "over the top" of Ethernet.

7.2 Architectural Elements

7.2.1 DPoE Network Definition

The term DPoE Network means all of the elements of a DPoE implementation, including at least one DPoE System and one or more DPoE ONUs connected to that DPoE System.

7.2.2 DPoE System Definition

The term DPoE System refers to that set of sub-systems within the hub site as illustrated in Figure 1, Figure 2, and Figure 3.

7.2.3 DPoE ONU Definition

The term DPoE ONU refers to the general case of an EPON ONU that minimally meets the DPoE specifications described in [DPoE-SP-MEFv1.0] as "Transport Mode." The term DPoE ONU refers to one individual device including at least the EPON ONU function, but also other embedded subsystems if implemented. There are two types of DPoE ONUs. These are the DPoE Standalone ONU and the DPoE Bridge ONU.

7.2.4 DPoE Standalone ONU Definition

The term DPoE Standalone ONU refers to a standalone DPoE ONU that operates without the need for any external device, such as a DEMARC device, and which provides the full functionality of all of the DPoE specifications.

A DPoE Standalone ONU MUST be an EPON ONU with no other function beyond [802.3ah] or [802.3av] and a minimum of one [802.3] Ethernet port. A DPoE Standalone MAY contain more than one Ethernet port and support a variety of [802.3] physical media. If a DPoE Standalone does have more than one Ethernet port, it MAY also support a variety of bridging and switching functions. A DPoE ONU MAY contain Ethernet ports inside of the device. Specifically, a DPoE Standalone ONU MAY include a PacketCable 2.0 eDVA, a DOCSIS 3.0 eRouter, eDSG, or other eDOCSIS implementation. Any eDOCSIS implementation must be accessible across the EPON by a dedicated LLID mapped to a logical Ethernet bridge or switched port as illustrated in Figure 1 and Figure 2. The LCI interface represents the [802.1] switch to eDOCSIS device interface providing an architecture equivalent to the eDOCSIS architecture. A DPoE Standalone ONU MUST use a dedicated LLID for every logical Ethernet bridge port or switched port to access any embedded Service Application/Functional Entity (eSAFE) connected to that bridged or switched Ethernet port, which is known as an LCI.

7.2.5 DPoE Bridge ONU Definition

The term DPoE Bridge ONU refers to a standalone EPON ONU that MUST have at least one S interface that can be operated in "Transport Mode" as described in [DPoE-SP-MEFv1.0]. One example of a DPoE Bridge ONU is an SFP-ONU. Another example of a DPoE Bridge ONU is a simple ONU based on an EPON chipset with one or possibly two Ethernet ports, but with no switching, forwarding, or encapsulation capabilities beyond the EPON chipset.

7.3 Interface and Reference Point Detailed Descriptions

DPoE specifications have both interfaces and reference points.

7.3.1 DPoE Interfaces

Interfaces are the inputs and outputs of the separable subsystems of DPoE Elements – that is, the inputs and outputs between the OSS and the DPoE System, between the DPoE System and a DPoE ONU, and between the DPoE ONU and various types of interfaces to CPE, eDOCSIS devices, and DEMARC devices. Since DPoE Standalone ONUs can contain embedded subsystems (eDOCSIS devices) that are often built by OEMs other than the ONU vendor, the DPoE specifications include a set of interfaces within the DPoE ONU called the S_1 interfaces.

7.3.1.1 D Interface

The D interface is the interface from the DPoE System to the operator's IP network for all IP traffic, including bearer traffic for IP service and OAMP functionality for IP and Ethernet services. The D interface on the DPoE System MUST support IP routing and forwarding as described in [DPoE-SP-IPNEv1.0]. The D interface is the interface for interoperability between the operator's OSS and DPoE Systems.

7.3.1.2 TU Interface

The TU interface is the interface between a DPoE System and one or more DPoE ONUs. The TU interface is the interface for interoperability between DPoE System and DPoE ONUs. Any DPoE System MUST interoperate with any DPoE ONU.

7.3.1.3 S Interface

The S interfaces and reference points, in general, represent the customer interfaces in the DPoE Network on DPoE ONUs.

7.3.1.4 S_1 Interface

S_1 interfaces are only on DPoE ONUs with external PHY interfaces operating as MI or MU. An S_1 interface on a DPoE Standalone ONU MAY be hardwired to an eDOCSIS device and be an LCI interface. An S_1 interface on a DPoE Standalone ONU MAY be configurable to act as an LCI on a port, as CMCI, MI, or MU interface. An S_1 interface on a DPoE Standalone ONU SHOULD be capable of being configured as one and only one of: LCI, CMCI, MI, or MU. An S_1 interface on a DPoE Standalone ONU MUST NOT be configured as more than one of these at any one time. Specifications for provisioning ports for MI or MU services are in [DPoE-SP-MEFv1.0]. Specifications for provisioning ports for CMCI or LCI service are in [DPoE-SP-MULPIv1.0].

7.3.1.5 LCI Interface

The LCI interface is equivalent to an eDOCSIS Logical CPE Interface (LCI) as described in [eDOCSIS]. An S_1 interface connected to an eDOCSIS device such as an eDVA, eRouter, or eDSG is always an LCI interface. An S_1 interface connected to an Ethernet port on a DPoE Standalone ONU MAY be configured as an LCI interface.

7.3.2 DPoE Reference Points

Reference points are similar to interfaces in that they point to locations where there are inputs and outputs. However, reference points may either not be separable because they are inside of an element, or they may be a reference to an independent industry standard.

7.3.2.1 MN Reference Point

The logical forwarding function that connects to the physical or logical interface that provides the MEF INNI to the operator's Metro Ethernet Network (MEN) is the MN reference point in DPoE. MN is a reference point because DPoE specifications do not specify the requirements for the interface. The requirements for MN are specified in MEF, [L2VPN], and IEEE specifications. The MN reference point is used to describe the MEF bearer traffic interface on a DPoE System connecting to an operator's MEN.

7.3.2.2 S Reference Point

The S interfaces and reference points, in general, represent the customer interfaces in the DPoE Network on DPoE ONUs.

7.3.2.3 S_2 Reference Point

The S_2 reference point is useful for the explanation of DEMARC device interfaces. Since the DEMARC device is not designed to DPoE specifications, the S_2 reference point is informative (only for reference). DPoE ONUs do not have access to port and interface information on DEMARC devices.

7.3.2.4 C_S Reference Point

The classifier on the DPoE System is identified and described by the Classifier-System (C_S) reference point. This reference point is useful for describing classification, scheduling, and forwarding functions required for interoperability between DPoE Systems and DPoE ONUs and for interoperability between DPoE Systems and the D and MN interface and reference points.

7.3.2.5 C_O Reference Point

The classifier on a DPoE ONU is identified and described by the Classifier-ONU (C_O) reference point. This reference point is useful for describing classification, scheduling, and forwarding functions required for interoperability between the OSS or DPoE System and DPoE ONU. It is also useful for describing the same parameters between the C_O reference point and the S interfaces or reference points.

7.3.2.6 MI Reference Point

The physical or logical interface that provides the MEF INNI to a DEMARC device (at a customer premise) is the MI reference point. MI is a reference point because DPoE specifications do not specify the requirements for the interface. The requirements for MI are specified in MEF and IEEE specifications. The MI reference point is used to describe the MEF bearer traffic interface on a DEMARC device, which in turn provides either another INNI (MI reference point) or MEF UNI interface (MU reference point).

7.3.2.7 MU Reference Point

An interface that provides the MEF UNI on a DPoE ONU or on a DEMARC device is the MU reference point. MU is a reference point because DPoE specifications do not specify the requirements for the interface. The requirements for MU are specified in MEF and IEEE specifications. The MU reference point is used to describe the MEF User to Network Interface (UNI) on DPoE Standalone ONUs or DEMARC devices.

7.4 Architectural Requirements

DPoE implementations are organized around an architecture designed for compatibility with:

- Existing DOCSIS Operations and Support Systems (OSS)
- Existing DOCSIS back-office systems
- Existing Layer 2 Ethernet, IP-based Ethernet Transport, or MEF MEN networks
- Existing EPON OLT systems and architecture
- Existing EPON ONUs and the SFP-ONU

All DPoE Network elements require [802.3ah], and, if applicable, [802.3av], [802.3], [802.1], and IETF protocol support (as defined in [DPoE-SP-IPNEv1.0] and [DPoE-SP-OSSIv1.0]). The architecture does not require any interoperability between sub-systems that comprise the DPoE System located in the operator's network facility. Each DPoE Network may be implemented in a variety of physical configurations. The DPoE architecture requires only the interfaces to the PON (the TU interface), the OSS and IP transport interfaces D, and the MEF Ethernet interfaces MN and MU to be interoperable with other systems.

7.4.1 Service Based Architecture

The service model for DPoE specifications is the Metro Ethernet Forum carrier Ethernet model using DOCSIS-based provisioning and operations. DOCSIS IP services are delivered "over the top" of virtual Ethernet services. Ethernet tags used for IP(HSD) service delivery cannot be exposed to DOCSIS IP(HSD) service customer interfaces (such as CMCI). In other words, the encapsulation and de-encapsulation for IP(HSD) [802.1ad] forwarding is performed in the DPoE ONU after the CMCI interface on ingress and before the CMCI interface on egress. Therefore, a DPoE ONU MUST NOT pass [802.1ad] frame headers for IP(HSD) forwarding past the CMCI interface. The MEF UNI interface for IP(HSD) services is exposed only to embedded subsystems (eDOCSIS devices) across the LCI interface. The LCI interface is a special case of an S interface that connects to an eDOCSIS device. IP(HSD) unicast services are built with E-LINE or E-Tree terminating at the DPoE System. Each [DPoE-SP-MULPIv1.0] IP Serving Group (IP-SG) is treated as an [802.1ad] service provider tagged network (S-VLAN). Each IP(HSD) customer is assigned a single [802.1ad] customer (C-VLAN). Other services may be built using Ethernet terminated elsewhere in the network where services interfaces are available.

The Metro Ethernet service model for external facing MEF UNI interface (to the MU interface) is specified in [DPoE-SP-MEFv1.0]. [DPoE-SP-MEFv1.0] is based upon the DOCSIS [L2VPN] provisioning model. However, [DPoE-SP-MEFv1.0] provides new encapsulation methods for use in DPoE Networks.

7.5 Service Model

DPoE services are based on the operator requirement to deliver a high quality of service for each service. DPoE services are built on the EPON LLID.

7.5.1 One to One Scheduler to Service Port Mapping

In order to provide strict scheduler control for each service, a DPoE ONU MUST support one or more dedicated LLIDs for each and every service. To provide strict control of the EVC traffic, management traffic SHOULD be separated from the EVC. This can be accomplished by providing a minimum of one (1) LLID for the management[3] of the entire DPoE ONU. Management within the EVC is permitted for embedded subsystems or external devices.

[3] Nothing in this requirement for the number of recommended LLIDs constitutes a requirement for IP-based management. DPoE ONUs are managed via OSSI through the DPoE System to the DPoE ONU with EPON OAM Extensions. See [DPoE-SP-OSSIv1.0] and [DPoE-SP-OAMv1.0] for those specifications which do not require a dedicated LLID.

However, each DPoE ONU S interface SHOULD support a dedicated LLID for management. This allows the operator to provision one LLID for the EVC path traffic and one LLID for OAMP to an external DEMARC device. This provides in-band management that is outside of the Metro Ethernet service EVC.

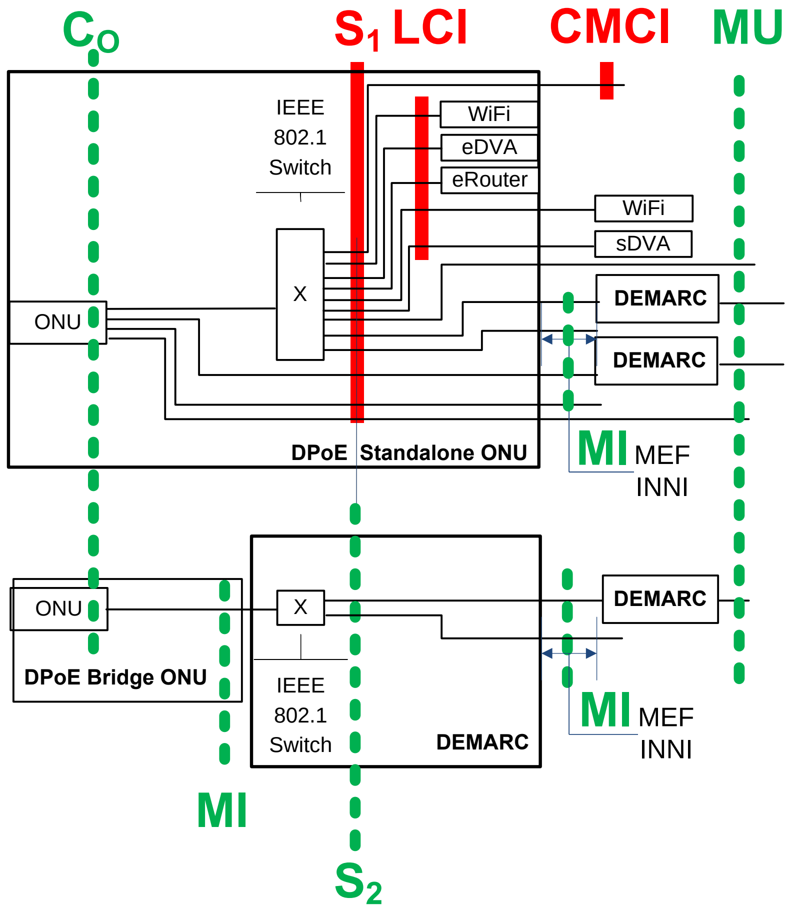

Figure 10 - DPoE ONU one to one classifies one service, from one port, to one LLID

All DPoE ONUs (applies to DPoE Standalone ONU and DPoE Bridge ONU) must comply with the following requirements:

A DPoE ONU S_1 interface MUST be configured and operated in one and only one of the following modes:

MEF

- MI interface
- MU interface

DOCSIS

- CMCI
- LCI interface

Specifications for provisioning ports for MI or MU services are in [DPoE-SP-MEFv1.0]. Specifications for provisioning ports for CMCI or LCI service are in [DPoE-SP-MULPIv1.0]. DPoE forwarding and configuration for

DOCSIS modes are described in [DPoE-SP-MULPIv1.0]. DPoE forwarding and configuration for Metro Ethernet services are described in [DPoE-SP-MEFv1.0].

A DPoE Bridge ONU only has a trunked Ethernet (Ethernet PHY, SGMII, XGMII, or other) interface. Although this specification does dictate requirements for DEMARC devices, the S_2 reference points on a DEMARC device SHOULD be counted as S interfaces to determine the LLID requirements or recommendations for a DPoE Bridge ONU in a particular implementation.[4] A DPoE Bridge ONU MUST support a minimum of 8 LLIDs.

An LCI interface MAY be a logical or physical Ethernet interface internal to a DPoE ONU. A DPoE LCI interface MAY be a physical Ethernet interface to an external or standalone DOCSIS device across a MEN at a customer premise. Examples of LCI interfaces include any possible eDOCSIS device, such as an eRouter, eDVA, eDSG, MoCA, sDVA, WiFi (internal or external), etc.

- Each eDOCSIS device is counted as a single LCI interface.
- Each LCI interface is a type of S_1 interface

"Ports" or "Interfaces" connected to embedded subsystems (such as the "air" interface on an WiFi access point, POTS or T.38 ports on an eDVA) SHOULD NOT be included in the port count for the minimum number of required ports. For example, an eDVA with two POTS or T.38 ports is counted as a single LCI interface, and therefore, as a single S_1 interface.

A DPoE ONU MUST support (at a minimum) the greater of:

1. an absolute minimum of eight (8) LLIDs per ONU; or
2. one LLID for each and every S interface on the DPoE ONU plus one dedicated for out-of-bearer (out-of-band) management

A DPoE ONU SHOULD support the sum of:

1. two LLIDs for each and every S interface on the DPoE ONU. One for customer (bearer) traffic and one for out-of-bearer (out-of-band) management traffic; plus
2. one LLID for optional IP or Ethernet based management[5] of the DPoE ONU.
(This could be used for applications like Performance Management, etc.)

A DPoE ONU MAY support the greater of:

1. two LLIDs for each eDOCSIS external port. For example, a four port eDVA, which is only a single LCI interface, and therefore, only a single S_1 interface, would have 4 LLIDs (one for each external POTS port) for the eDVA; plus
2. two LLIDs for each eDOCSIS interface that is a single port S_1 interface (such as an LCI interface used for IP service); plus
3. two LLIDs for each S interface operating as an MI interface; plus
4. two LLIDs for each MU interface on a DEMARC device connected via an MI interface; plus
5. one LLID for management of the DPoE ONU.

In summary, the LLID requirements are as follows:

S is the number of S (including S1 + S2) interfaces and references points on the DPoE ONU

---

[4] Counting each DEMARC device port (S2 reference point) as an S interface in determining the LLID requirements results in a set of LLID requirements (MUST, SHOULD, and MAY) that allow for one dedicated LLID for each DEMARC device port plus one LLID for managing the DEMARC device (MUST), and additional recommendations for management of services or functions for each individual service on each port (SHOULD and MAY).

[5] Nothing in this requirement for the number of recommended LLIDs constitutes as requirement for IP-based management. DPoE ONUs are managed via OSSI through the DPoE System to the DPoE ONU with EPON OAM Extensions. See [DPoE-SP-OSSIv1.0] and [DPoE-SP-OAMv1.0] for those specifications which do not require a dedicated LLID.

LCI is the number of LCI interface on the DPoE ONU

Minimum number of LLIDs a DPoE ONU MUST have = $\min_8 〚(S+1)〛$

Number of LLIDs a DPoE ONU SHOULD have = $\min_8 〚((S \times 2)+ 1)〛$

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8[(S-LCI)x\ 2+(LCI \times 2)+ 1]$

Since the only difference between the "SHOULD" and "MAY" requirements is for LCI interfaces, a DPoE ONU without LCI interfaces will always have the same "SHOULD" and "MAY" requirements. Please refer to Appendix IV for examples of DPoE ONU LLID requirements.

7.5.1.1 Single Tenant or Single Customer

Customers that require four services such as: Ethernet, VoIP (whether on an eDVA or externally on an ATA), IP-TV, and IP, would require at least 5 LLIDs. Such an implementation may require additional LLIDs if an external system, such as a VoIP ATA, requires a dedicated management circuit for management, external to the DPoE System. Such a dedicated LLID could, for example, be mapped to a VLAN or other Ethernet tagged circuit and run on private [RFC 1918] address space for external management by OSS elements out of scope of [OSSIv3.0] and [DPoE-SP-OSSIv1.0].

7.5.1.2 Multiple Tenant or Multiple Customer

Multiple Tenant Units (MTU) or other locations with multiple customers could best be served with either a multiple port DPoE Standalone ONU or a DPoE Bridge ONU connected to a DEMARC device.

7.5.1.3 External Devices or CPE

The DPoE ONU is not a CPE. The term CPE is reserved for use for the customer's equipment as specified in [DOCSIS].

DPoE specifications do not address the provisioning of separate circuits for external device management. Such circuits may be provisioned by operators as EVCs within an EPL according to the specifications of [DPoE-SP-MEFv1.0]. An operator may also provision VLANs [802.1q] or [802.1ad] (S-VLAN or C-VLAN) tags for in-band management (of a DEMARC device) outside of the Metro Ethernet service EVCs.

7.5.2 Many to One

DPoE services MAY use more than one LLID for a service. An operator could offer services that support managed QoS of EVCs within an EPL, E-VPL, or E-LAN. Similarly, an operator could offer services that support managed QoS within a DOCSIS IP service.

7.5.3 Any to Any

DPoE ONUs SHOULD NOT permit the classification of frames or packets from a port provisioned for one service to the LLID mapped to a port provisioned for a different service. The specifications for forwarding and classification are in [DPoE-SP-MULPIv1.0] and [DPoE-SP-MEFv1.0].

7.5.4 NNI Interfaces

The DOCSIS IP(HSD) NNI is referred to as the D interface on the DPoE System and the LCI interface(s) on the DPoE ONU. The Metro Ethernet service NNI and UNI interfaces (or reference points) are described as the MN and MU interfaces or references points on the DPoE System and DPoE ONU, respectively.

Classification and scheduling in a DPoE Network is defined by the DPoE System classification reference point C_S and the DPoE ONU classification reference point C_O.

7.5.4.1 Upstream Classification Reference Point

DPoE ONUs MUST perform all upstream traffic encapsulation, tagging, marking, and transcoding of frames or packets into the format expected for EPON classification at or before the C_O reference point. Upstream traffic encapsulation and forwarding are described in [DPoE-SP-MULPIv1.0] for IP(HSD) and [DPoE-SP-MEFv1.0] for Metro Ethernet services. Classification and scheduling for all services are described in [DPoE-SP-MULPIv1.0]. By definition, all upstream classification is performed at the C_O reference point. Some traffic encapsulation, tagging, marking, and transcoding of frames or packets is performed on the DPoE System after the C_O reference point, at the C_S reference point, or even after the C_S reference point, but always prior to the D or MN interface.

7.5.4.2 Downstream Classification Reference Point

DPoE Systems MUST perform all traffic encapsulation, tagging, marking, and transcoding at or before the C_S reference point. Downstream traffic encapsulation and forwarding are described in [DPoE-SP-MULPIv1.0] for IP(HSD) and [DPoE-SP-MEFv1.0] for Metro Ethernet services. Classification and scheduling for all services are described in [DPoE-SP-MULPIv1.0]. By definition, all downstream classification is performed at the C_S reference point.

7.6 Service and Device Management Architecture

DPoE ONUs are managed primarily using the EPON OAM [802.3ah] and OAM Extension in [DPoE-SP-OAMv1.0].

DPoE ONUs are managed from OSSI by proxy through a virtual Cable Modem (vCM) as described in [DPoE-SP-OSSIv1.0].

Embedded systems, including any eDOCSIS devices, are operated "over the top" of EPON by the DPoE System using [802.1] tagged Ethernet frames for traffic isolation that can be readily mapped to SF (and therefore dedicated LLIDs). This may be accomplished by using DOCSIS IP classifiers (as specified in [DPoE-SP-MULPIv1.0]) to classify CMCI or eDOCSIS traffic to an SF (and therefore a dedicated LLID). DPoE Standalone ONUs operate IP(HSD) forwarding for eDOCSIS LCI or CMCI interfaces using this same single model. When a DPoE Standalone ONU is (optionally) an eDOCSIS device, any S1 interface not directly connected to an eSAFE may be an LCI (to an external or standalone SAFE) or a CMCI port. When a DPoE Standalone ONU is not an eDOCSIS device, the DPoE Standalone ONU performs the same traffic encapsulation, classification, and forwarding on a per S1 'interface' basis. Whether the S1 'interfaces' are logically terminated at an eSAFE, treated as an LCI or as a CMCI (as in a DPoE Standalone ONU without eDOCSIS), the forwarding is the same for all IP(HSD) configurations. The following specifications must be followed for IP(HSD) service and all S1 reference points (which are LCI interfaces on an eDOCSIS device):

- The DPoE System MUST use the [802.1ad] tags to transmit IP(HSD) SF downstream on the TU interface.
- The DPoE System MUST drop the [802.1ad] tags from IP(HSD) SF prior to the D interface on the upstream.
- The DPoE Standalone ONU MUST use the [802.1ad] tags to transmit IP(HSD) SF upstream on the TU interface.
- The DPoE Standalone ONU MUST drop the [802.1ad] tags prior to transmitting IP(HSD) downstream at the S1 reference point

The D interface MAY use Ethernet tags for other purposes as described in [DPoE-SP-IPNEv1.0]. Such tags will not conflict with tags used in DPoE services if the above requirements are met.

External systems connected to a DPoE ONU across an S interface are managed across a provisioned Ethernet service as specified in [DPoE-SP-MEFv1.0] or by using an IP path across the LCI interface if DOCSIS IP-based management is acceptable. External systems connected to a DPoE ONU across an S interface MAY be managed within a single EVC providing both traffic and management services to the customer. External systems connected to a DPoE ONU across an S interface MAY also be managed within a single EVC providing only management services for the operator to remotely operate a customer's CPE separate from an EVC for the customers traffic.

7.7 IP Forwarding Model

The IP forwarding model is described in detail in [DPoE-SP-MULPIv1.0]. The DPoE System MUST include a fully functional IP router to provide the functions described in [DPoE-SP-MULPIv1.0] and meet all of the system requirements in [DPoE-SP-IPNEv1.0]. These are the same requirements that operators have for DOCSIS-based CMTS systems.

7.7.1 MEF Model

The Metro Ethernet Forum architecture (MEF 6) and model for service is a stacked VLAN model that can be implemented with [802.1ad]. There are three main differences between the Metro Ethernet services requirements in [DPoE-SP-MEFv1.0] and the IP(HSD) requirements in [DPoE-SP-MULPIv1.0]. First, [DPoE-SP-MEFv1.0] offers a variety of encapsulation modes (to support legacy equipment and optionally PBB). Second, some QoS parameters for MEF differ from those in DOCSIS IP(HSD) service. Third, the provisioning of Ethernet transport in [DPoE-SP-MEFv1.0] is explicit in the L2PVN portion of the CM configuration file but is local to the DPoE System and fully autonomous for IP(HSD) DOCSIS service as specified in [DPoE-SP-MULPIv1.0]. The IP(HSD) model does not require provisioning of Metro Ethernet service by the DOCSIS OSS. The provisioning of MEF forwarding is done locally on the DPoE Network by the DPoE System.

7.7.2 MEF Rationale for IP(HSD)

The MEF model is a Carrier Ethernet model. This provides flexibility for hiding layers of VLAN complexity (in the transport) from the IP network. The MEF model implemented with [802.1ad] offers 4096 VLANs for VLAN Identifier (VID). With S-VLAN and C-VLAN, there are (4096 x 4096) 16,777,216 unique S-VLAN + C-VLAN pairs possible. This means that a single MEN could support up to that number of unique ports/services. If the S-VLAN and C-VLAN are used only on the DPoE Network (and not extended to the MEN), that number of LCI ports could be made available for each DPoE System. A single [802.1q] (Q tag) based system would support only 4096 ports. That is not enough for a DPoE System. [802.1q] VIDs would also overlap with S-VLAN VIDs in a mixed environment with both MEF and IP(HSD) services if IP(HSD) relied on "simple" Q tags.

MEF forwarding for IP(HSD) uses the C-VID as an "inner tag" and the S-VID as the outer tag without respect to the typical name convention. The "Customer" or "C" VID is, for DPoE IP(HSD), just an inner tag that carries the IP(HSD) service for a single eSAFE or CMCI.

7.7.3 DPoE IP(HSD)

[DPoE-SP-MULPIv1.0] implements eDOCSIS using [802.1] bridge running IEEE 802 S, C, and D components[6] to provide an LCI interface to each eSAFE. The MEF model describes exactly such a model with S, C (as optional) and D components and calls each virtual connection an EVC. In summary, the DPoE model to support LCI is an EVC for each eSAFE.

[6] An "S component" is an S-VLAN forwarding implementation on an IEEE 802.1 switch. A "C component" is a C-VLAN forwarding implementation on an IEEE 802.1 switch (sometimes called a VLAN). A "D component" is a [802.1d] or MAC bridge forwarding implementation on an IEEE 802.1 switch (sometimes called a bridge group).

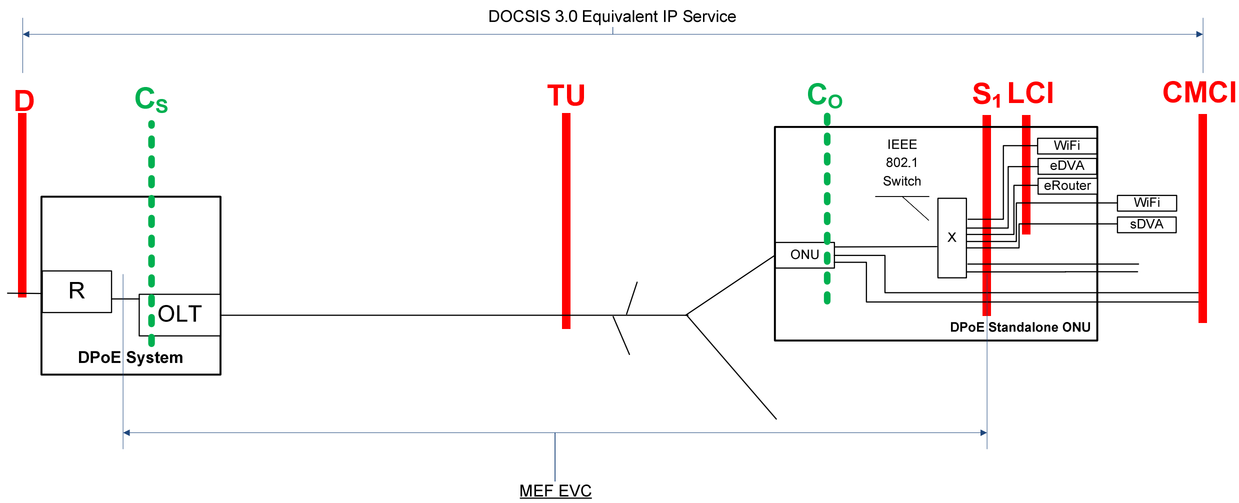

Figure 11 - DPoE IP(HSD) high level view

7.7.4 DPoE System Future Extensibility

The MEF EVC based architecture for IP(HSD) allows for service, address, or other S-VLAN based grouping and trunking of traffic from an OLT to the IP router or other system components. While the provisioning of the system would remain monolithic, the forwarding model could be more easily separated. Although there is no "modular" DPoE architecture, the S-VLAN trunking could be used within a single vendor to build multi-shelf systems.

Similarly, the S-VLAN trunking could be used for MEN transport within the operator's network. In other words, the DPoE System could be implemented using MEN between or even within the DPoE Elements defined in Section 3.

In the customer premise, the use of Metro Ethernet services also allows for the future possibility of a MEF INNI to transport EVCs from a DPoE ONU across a MEN at the customer premise to a standalone DOCSIS device that is not directly attached as [eDOCSIS] currently assumes. Such an implementation could, for example, use an sDVA as a DPoE ONU to accomplish the same functions and achieve the same service provisioning and automation as a DPoE Standalone ONU with an eDVA.

7.7.5 IP(HSD) Implementation of MEF with [802.1ad]

The implementation of the forwarding model within the DPoE System is not specified in DPoE. As long as the DPoE Network acts as expected and can be tested at the D and TU interfaces, the implementation within the DPoE System is vendor-specific. At the TU interface, the IP(HSD) traffic over EPON traffic can be inspected for [802.1ad] encapsulation.

Figure 12 below shows a logical view of a DPoE System with two IP router interfaces, DR_1 and DR_2. These are the default routers for two different IP subnets operating on the DPoE System. Each subnet can be configured by operators for their own policy and functional needs. In this example, the IP subnet 1.1.1.0/24 is configured to provide service for all CMCI and eRouters. Each CMCI interface and eRouter interface is connected by an EVC to the Default Router (DR) on that subnet DR_1 at 1.1.1.254 (as in Table 3). The IP subnet 2.2.2.0/24 is configured to provide service for voice service (including all eDVA and sDVA). Each DVA interface is connected to an LCI which is connected to an EVC to the Default Router on that subnet, DR_2, at 2.2.2.254 (as in Table 3).

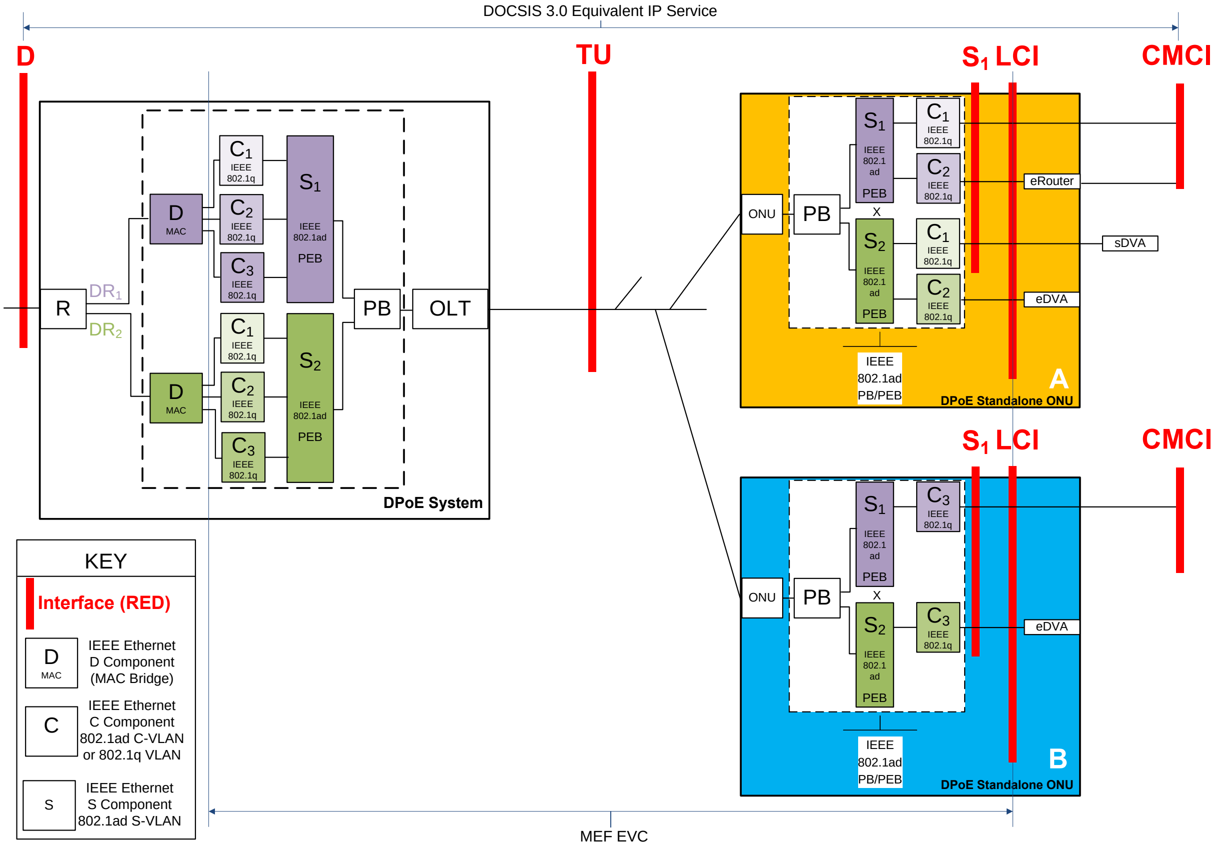

Figure 12 - IP(HSD) Implementation of MEF with [802.1ad]

In the example of Figure 12, S-VLAN ID = 1 (S_1) is configured by the operator to provide IP service for all eRouter type eSAFEs and all CMCI interfaces. S-VLAN ID = 2 (S_2) is configured by the operator to provide IP service for all DVA type eSAFEs (both embedded and standalone).

Note that each S-VLAN ID is available on any DPoE ONU connected to a single PON port on a DPoE System. Thus, each S-VLAN ID is available across the DPoE Network.

Table 3 - IP(HSD) EVCs

Default Router (DR) Name	Interface Type	eSAFE Type	S-VLAN ID	C-VLAN ID
DR_1	CMCI	na	1	1
DR_1	LCI	eRouter	1	2
DR_2	LCI	sDVA	2	1
DR_2	LCI	eDVA	2	2
DR_1	CMCI	na	1	3
DR_2	LCI	eDVA	2	3

As configured in this example, each S-VLAN corresponds to an IP subnet. This is not necessary, but just one example. One IP subnet could be served by multiple S-VLANs. Likewise, S-VLANs may service multiple IP subnets. There is no necessary relationship between the two. Normally this is not possible because an IP subnet is usually associated with a broadcast domain (a VLAN).

Implementing the system this way has two effects. First, it eliminates the need for MAC address. When VLANs are used as point-to-point circuits, there is no longer a need to use the Layer 2 (MAC) address for forwarding. Each S-VLAN + C-VLAN combination forms a unique point-to-point circuit (like a PVC in ATM). Like other point-to-point technologies, a MAC address is not required. This is described in further detail in Appendix III.5 below. Second, the traffic from the DR (to a given destination IP address) can be forwarded to one and only one S-VLAN ID + C-VLAN ID combination. Likewise, traffic from an S-VLAN + C-VLAN combination can be forwarded only to the DR on one side, and to the CMCI or LCI on the other. When configured to operate in this manner, it is not possible for traffic to traverse from one S-VLAN ID + C-VLAN ID pair to another. Another way to think about this is that each S-VLAN + C-VLAN combination is a two-port bridge group. The DR has multiple logical interfaces with one in each of the S-VLAN + C-VLAN bridge groups. The other port on the bridge group connects to one and only one LCI or CMCI interface. In short, there is no broadcast domain. For Carrier Ethernet and operators, this provides a substantial benefit. With broadcast networks eliminated, the access network is far more secure. It is not possible for a MAC address to be spoofed.

7.8 eDOCSIS

DOCSIS 1.1 operates what we now call a "bridged" DOCSIS service at the Cable Modem CPE Interface (CMCI). DOCSIS 3.0 supports a bridged as well as a routed option utilizing an eRouter with embedded IP routing on the CM. IP services that operate on the DOCSIS model do so across the eDOCSIS Logical CPE Interface (LCI).

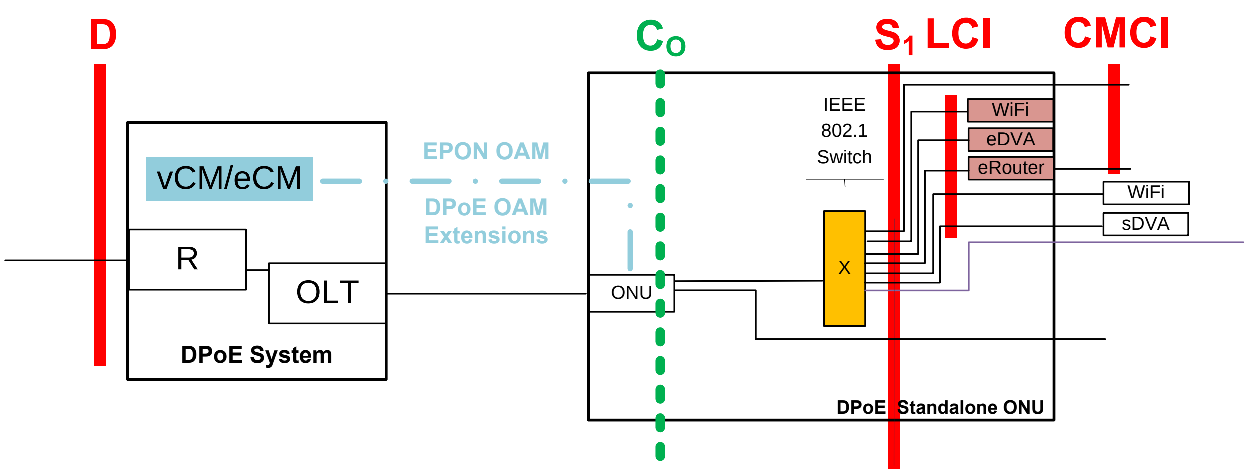

Figure 13 - DPoE eDOCSIS Reference Model

7.8.1 eDOCSIS on DPoE Standalone ONU

Any DPoE Standalone ONU that implements a DOCSIS eSAFE such as eRouter, eDSG, or eDVA is an eDOCSIS device and MUST support [eDOCSIS].

A DPoE Standalone ONU that is (optionally) an eDOCSIS device MUST have an IP stack on the DPoE ONU for each eSAFE. This includes eDVA, eRouter, eDSG, WiFi, or any other eSAFE subsystem within a DPoE Standalone ONU. Standalone DOCSIS devices such as an sDVA use IP stacks on each respective device. The vCM model is only used for DOCSIS CM management functions.

A DPoE Standalone ONU that is (optionally) an eDOCSIS device MUST implement the "SF Classifier" and "MAC Bridge and Filter" functions of the eCM.

A DPoE Standalone ONU MAY offer "Physical CPE interfaces" that provides CMCI or other functionality. A DPoE Standalone ONU MAY offer any "Physical CPE interfaces" that support Ethernet.

7.8.2 eDOCSIS Proxy on DPoE Standalone ONU

Any DPoE Standalone ONU that does not have any eSAFEs MAY support external or standalone DOCSIS devices such as sDVA. A DPoE Standalone ONU MAY offer a means to proxy sDVA, DEMARC, or other device capabilities and information in order for the DPoE Standalone ONU to have the necessary data to forward to the DPoE System for discovery (as described below) and act as an eCM.

7.8.3 eSAFE Discovery

A DPoE Standalone ONU that is an eDOCSIS device (as defined above) MUST support the method described in [DPoE-SP-MULPIv1.0] for device capability discovery. The device capability discovery is a proxy for the parameters required for DHCP options to identify eSAFEs as specified in [eDOCSIS]. In order to identify external devices as eSAFEs, the DPoE System MUST support the method described in [DPoE-SP-MULPIv1.0] for device capability discovery.

7.8.4 vCM and eCM

DPoE Systems MUST virtually implement part of the eDOCSIS eCM within the DPoE System as illustrated in Figure 7. The vCM in the DPoE System MUST operate as the eDOCSIS eCM for management purposes. The DPoE System can use the DHCP Options, detailed in [eDOCSIS], to identify if eSAFEs are present in a DPoE Standalone ONU. When a DPoE System receives and proxies forward a software image for a DPoE Standalone ONU that is an eDOCSIS device, the DPoE Standalone ONU software image MUST be a monolithic software image for the eDOCSIS device, which is a DPoE Standalone ONU with eSAFEs, as described in [eDOCSIS].

DPoE Systems MUST use [DPoE-SP-OAMv1.0] PDUs for all required eCM functions as described in [eDOCSIS]. DPoE Standalone ONUs MUST use [DPoE-SP-OAMv1.0] PDUs for all required eCM functions as described in [eDOCSIS].

7.8.5 LCI Extended

eDOCSIS allows for "other eSAFEs" with "other interfaces" (as described in [eDOCSIS]). DPoE specifications support "other eSAFEs" within the eDOCSIS device or with external interfaces. An example might be a MoCA or WiFi device which acts as a PHY bridge, allowing other external devices or systems to be served by an LCI. eDOCSIS also specifically includes an interface from the bridge directly to a PHY ("Physical CPE interfaces") to "other CPEs." [DPoE-SP-MULPIv1.0] allows for this same variation but explicitly treats that connection as just another LCI.

With every connection treated as an LCI in DPoE Standalone ONU (as shown in Figure 14), there is a uniform model for both embedded and standalone subsystems. An operator could choose an embedded PacketCable 2.0 Digital Voice Adapter (eDVA) or a standalone PacketCable 2.0 Digital Voice Adapter (sDVA), and the provisioning and operation of the DPoE Standalone ONU will be the same. Combined with a bridge interface like MoCA, it would be possible to support eDSG over MoCA with the LCI extended over MoCA.

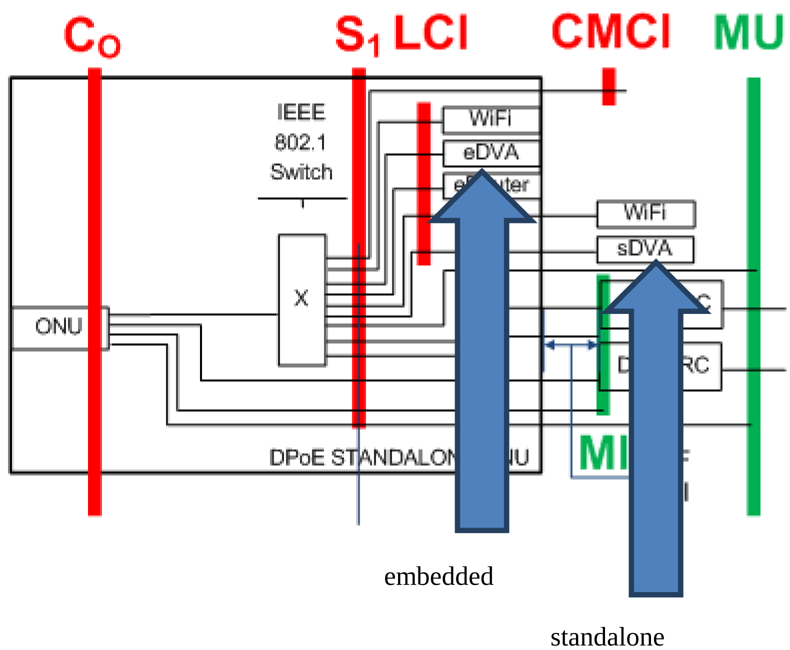

Figure 14 - DPoE Specifications support both embedded and standalone SAFEs

Appendix I Introduction to Multi-Access Optical Networks

Most optical network technologies in use today allow only for a single optical transmitter and a single optical receiver on each strand of glass. This is called simplex operation. Examples of simplex optical network technologies include 1000Base-SX, 1000Base-LH, 100BaseFX. Baseband technologies operate with only a single transmission channel across the medium (optical fiber). Most broadband optical technologies also use simplex operation. For example, most coarse wavelength division multiplexing (CWDM) and dense wavelength division multiplexing (DWDM) systems are composed of paired transmitters and receivers. Although wavelength division multiplexing (WDM) systems multiplex multiple signals onto a single strand of glass, these systems function by demultiplexing the signals at the receive side and sending each wavelength to a separate receiver. Each set of transmitters, multiplexers, fiber, demultiplexer, and receiver is simplex. Duplex communications are achieved by providing a second and separate set of transmitters, multiplexers, fiber, demultiplexer, and receiver for traffic in the opposite direction. Complex variations on DWDM, like reconfigurable optical add-drop multiplexing (ROADM) technology, are sophisticated multiplexers that can selectively add and drop wavelengths. Since these are fixed (provisioned) models, they are not multi-access optical networks. Like the baseband equivalents, there is always a one to one relationship between a transmitter and a receiver. Although the technology appears to offer more than single span topologies (such as rings), these logic constructs are not operated by the network in real time. ROADMs provide an optical domain control for setting up point-to-point circuits. Most DWDM systems are circuit-based models. There are duplex optical fiber technologies other than PON. [802.3ah] specified 100Base-BX10, which is a bi-directional (full duplex) standard that uses CWDM to multiplex two transmissions (both directions) onto a single strand of glass for 100Mbps Ethernet. The SCTE 55-1/55-2 standards specific the Hybrid-Fiber Coaxial (HFC) architecture, which supports duplex operation over the optical (fiber) portion of the network using WDM for transmission separation for each direction.

EPON allows simultaneous upstream and downstream transmission by using different frequencies in each direction. In the downstream, only the OLT transmits. With only a single transmitter, the timing of transmissions is managed within the OLT, without the need to schedule transmissions with other devices. Upstream transmission in EPON is scheduled across multiple ONUs. To accomplish this, EPON implements Time Division Multiple Access (TDMA) to separate transmission from each ONU over time. PON is the most widely used multi-access optical network technology currently available.

PONs differ from all of the simplex technologies because they allow multiple transmitters to share a single strand of fiber. Transmissions on PON are duplex. This means that transmission in both directions can occur simultaneously. This is accomplished by combining optical transmissions together onto a single path.

A PON is a particular type of multi-access optical network called a point to multipoint (P2MP) network. It is a type of graph called a tree. The unique property of a tree is that there is always and only one path through a set of required vertices to get from one edge (end point) of the tree to another edge (end point) of the tree. A PON is more specifically a rooted tree. The root of the tree is an edge that is called an Optical Line Terminal (OLT). The vertices in the tree of a PON are passive optical splitter combiners (OSC). The vertices allow paths from any edge to the root edge and from the root edge to any other edge. In a PON, there is no path from an edge (other than the root) to any other edge. These are important characteristics of a PON because they both permit and limit the flows of traffic (the paths) that can be used for unicast, broadcast, and multicast traffic. The implications for analog broadcast, Ethernet, IP, IP multicast, and security are very significant. These attributes of PON allow for efficient broadcast (and multicast) and significantly reduce the security needs and vulnerability of PON technologies. The passive attribute of PONs is also an important factor in the reduced operating cost of an ODN.

Appendix II Introduction to EPON

II.1 Downstream Channel

In the downstream channel, data frames broadcast by the OLT pass through a 1xN OSC or an OSC cascade to reach the connected ONUs. Each ONU receives a copy of every downstream data packet. The number of connected ONUs varies depending on the operator requirements. The number of connected ONUs is limited only by the available optical power budget and the bandwidth demand distribution among the connected end-users. A single 1G-EPON OLT could support as many as 2^{15}-1 LLIDs (the LLID field has the effective width of 15 bits). In 10G-EPON, an OLT can only support 2^{15}-255 because a block of 255 LLIDs was reserved for support of broadcast services in the future.

The downstream channel properties in this EPON system make it a shared-medium network. Packets broadcast by the OLT are selectively extracted by the destination ONU, which applies simple packet-filtering rules based on MAC and LLID.

II.2 Upstream Channel

In the upstream direction, EPON operates in the Point-To-MultiPoint mode (P2MP). All ONUs transmit their data packets to a single receiver module located in the OLT. Since the OSC is a directional device, an ONU cannot see the signal transmitted upstream by any other ONU. The appearance of P2P connection is achieved by centrally managed access to the upstream channel. The OLT allows only a single ONU at a time to deliver pending packets.

Since all ONUs belong to a single collision domain, centrally managed channel access is required. ONUs are not allowed to transmit any data unless granted specifically by the OLT. Data collisions are avoided because the OLT is aware of the scheduled transmissions from individual ONUs. The only exception from this centrally managed upstream channel access scheme is the so-called Discovery Process (as defined in [802.3ah] / [802.3av]), where new and not initialized ONUs are allowed to register in the EPON system.

II.3 Physical Layer

EPON system specifications can be divided into two generations. 1G-EPON specifications were released in [802.3ah], while higher speed 10G-EPON specifications were released in [802.3av]. All 1G-EPON related specifications are described in Clause 60 (PHY), 64 (MPCP), and 65 (Extensions to RS and PCS sublayers) – several references ALSO exist in other subclauses of [802.3]. 10G-EPON specifications will remain a stand-alone document (Clause 75 (PHY), Clause 76 (PCS), and Clause 77 (MPCP)) for the time being, expected to be incorporated into the base document of the IEEE Std. 802.3 standard.

1G-EPON [802.3ah] supports two different PHYs with a Channel Insertion Loss (ChIL) equal to or smaller than 20 dB (the so-called PX10 PHYs) and equal to or smaller than 24 dB (the so-called PX20 PHYs). All the parameters describing the PX10 and PX20 PHYs (both normative and informative) are included in Clause 60 [802.3ah]. Both PHYs support normatively the split ratio of at least 1:16 at the distance of 10 and 20 km, respectively. Detailed information on the particular parameters of the compliant transmitters and receivers for both types of PHYs can be found in Table 60-3 and 60-5 for Tx and Rx units in PX10 links and Table 60-6 and 60-8 for Tx and Rx units in PX20 links.

There is a new PHY (PX30), which was not included in [802.3ah] specifications. PX30 supports a ChIL of 29 dB and is used to operate on a connectorized fiber plant without fixed splices.

10G-EPON [802.3av] introduces the idea of an asymmetric data rate EPON that operates at the data rate of 10 Gbit/s downstream and 1 Gbit/s upstream (all PMDs of PRX type). Additionally, symmetric data rate EPON is supported, which operates at the data rate of 10 Gbit/s in downstream and upstream (all PMDs of PR type). Most of the power budget classes were devised focusing on backward compatibility with [802.3ah]. EPON (PR10/PRX10) is backward compatible with PX10. PR20 and PRX20 are backwards compatible with PX20. Issues related to the coexistence of [802.3ah] equipment with [802.3av] equipment on the same PON plant were foremost in the development of [802.3av].

II.4 Data Link

All data exchanged between the MAC entities located in the OLT and ONUs is transmitted as MAC datagrams (frames) in IEEE 802 Ethernet format.

EPON systems do not support frame fragmentation. A frame constitutes the smallest available unit of data that can be exchanged between two MAC entities over Ethernet medium (PON or other P2P medium). The MAC control signaling is transmitted in-band in the data channel with the information encapsulated into standard Ethernet frames relayed with the highest possible priority. Provided that higher layers need to transmit a chunk of data with the size exceeding the IEEE 802 standard 1500 byte Ethernet frame payload, such layers must rely on mechanisms (transport protocols) different from Ethernet for the fragmentation on the transmitter side and reassembly on the receiver side.

1G-EPON PX10 and PX20 PHYs are symmetric in terms of the data rate, i.e., the OLT PHY transmits at the nominal data rate of 1.25 GBd downstream and receives the 1.25 GBd data stream from the connected ONUs in the burst mode fashion. The effective data rate available for the customer data is, however, lower and equal to 1.0 Gbps (at the PCS service interface – see Figure 60-1 in [802.3ah]).

The 25% increase in the data rate is the direct result of the channel encoding scheme used in the 1000BASE-X connections (1G-EPONs use the 1000BASE-X MAC with P2MP specific extensions at the RS, PCS, PMA and MAC Client layers), i.e., the 8b/10b, which encodes 8 bits (a word) from the PCS service interface into 10 bits at the PMA service interface. Such an encoding helps achieve DC-balance (equal number of 1s and 0s in the resulting bit stream as transmitted in the given medium) and bounded disparity, and yet provide enough state changes to allow reasonable clock recovery. This means that the difference between the count of 1s and 0s in a string of at least 20 bits is no more than 2, and that there are not more than five 1s or 0s in a row (as defined in [802.3], clause 36.2.4.3). This helps to reduce the demand for the lower bandwidth limit of the channel necessary to transfer the signal, therefore providing savings on the electronic components required for both transmission and reception of the transmitted signal.

10G-EPON PHYs are divided into two groups in the function of the upstream data rate, i.e., (a) asymmetric data rate (PRX type) and (b) symmetric data rate (PR type). The PRX type PHYs operate at the data rate of 10.3125 GBd in downstream and 1.25 GBd upstream (effective data rate of 10 Gbps and 1 Gbps, respectively), using 1G-EPON PHY specifications for the upstream transmission channel and 10G-EPON PHY specifications for the downstream channel. The PR type PHYs operate at the data rate of 10.3125 GBd in downstream and upstream (effective data rate of 10 Gbps, using only 10G-EPON PHY specifications.

The 3.125% increase in the data rate is the direct result of the channel encoding scheme used in the 10G Ethernet links (10G-EPONs use the 10GBASE MAC with P2MP specific extensions at the RS, PCS, PMA and MAC Client layers) 64b/66b. 64b/66b encodes 64 bits (a word) from the PCS service interface into 66 bits at the PMA service interface. The channel encoding in 10G-EPON does not condition the signal in the same was in 1G-EPON; hence the overhead is much lower and the resulting transmission overhead limited.

II.5 FEC

The Forward Error Correction (FEC) is used in general to enhance the available power budget by increasing the receiver sensitivity and supporting the lower Bit Error Ratio (BER) levels. When the FEC is enabled, the transmitter adds redundant information to transmitted data stream, also known as an error-correction code, which increases resilience to bit errors. This allows the receiver to detect and correct errors (within some bounds) without the need to ask the transmitter to resend the data at the Ethernet layer. Of course, higher layer protocols can (optionally) implement additional data error detection, correction, or retransmission.

II.6 Stream-Based FEC Versus Frame-Based FEC

1G-EPON adopted optional frame-based FEC, while 10G-EPON uses mandatory stream-based FEC. Both mechanisms can provide extended protection against bit errors occurring during transmission in the optical channel. However, both are also quite different in many ways.

A stream-based FEC mechanism processes Ethernet frames and IDLEs as a stream of data symbols, resulting in a much simpler implementation, which is critical for high data-rate systems. This particular FEC encoding method requires both transmitter and receiver communicating over a physical medium to use the very same framing structure. A device not supporting FEC encoding will not be able to retrieve data and separate it from parity. This means that all ONUs in a 10G-EPON must use FEC. In the stream-based method, the parity symbols generated after each data block are inserted immediately after the FEC parity codeword that they are protecting, resulting in an interleaving pattern of data blocks and parity blocks.

In the frame-based method, the parity symbols generated for each block are grouped together and are appended at the end of a frame. This leaves the data frame itself unaltered, representing a major advantage of this particular encoding method. Any device not supporting FEC encoding may still receive the data, though will not take advantage of the enhanced FEC bit protection. In 1G-EPON, adoption of this particular FEC coding method allows for mixing ONUs with enabled and disabled FEC on the same ODN.

II.7 Upstream Channel and Burst-Mode

An EPON upstream channel uses a burst mode (similar to all P2MP systems with TDMA in the upstream channel), where the transmissions (slots) from multiple ONUs are time multiplexed and have a dead zone (guard band) between them to allow for turning the laser off and on between subsequent transmissions, adjust power levels and align to the incoming data stream and retrieve clock (CDR).

Due to possibly unequal distances between the OLT and the ONUs, optical signal attenuation in the PON may differ slightly for each connected ONU. The power level of the ONU transmissions in the upstream channel may also vary from one time slot to another (near-far problem). In order to detect the information carried in the incoming bit stream, the OLT receiver must adjust its decision threshold at the beginning of each received burst, which is commonly referred to as Automatic Gain Control (AGC). This reception mode, in which an optical signal arrives at the receiver in bursts with varying power levels, is called burst-mode reception.

In addition to performing the AGC function on the incoming data stream, burst-mode receivers must be able to acquire phase and frequency lock on an incoming signal, commonly referred to as clock and data recovery (CDR) functionality. The ability to perform AGC and CDR very quickly is paramount for a receiver to operate in burst mode, which is required fortunately only at the OLT (thus allowing for more complex and optimized solutions). The ONUs receive a continuous bit stream (data or idle signal) sent by the OLT over a P2P data link emulated on top of the P2MP PON structure, and thus do not need to readjust the receiver gain between individual frames. In fact, gain adjustment in the ONU receivers is carried out rather rarely, if at all. Typically, the ONU receiver dynamic range alone is sufficient to compensate for slight changes in the OLT transmitted power in the downstream channel.

In a TDMA PON, it is not enough to prevent ONUs from sending any data between the assigned timeslots, since, even in the absence of data being delivered to the laser driver by the MAC stack, lasers generate spontaneous emission noise. Spontaneous emission noise from several ONUs located close to the OLT can easily obscure the signal from a distant ONU (capture effect). ONUs must shut down their lasers at the end of each transmitted burst. The mode of operation in which the laser is being completely turned off between the transmissions is called burst-mode transmission. Because a laser cools down when it is turned off, and warms up when it is turned on, its emitted power may fluctuate at the beginning of a transmission, and thus this part of the upstream transmission slot cannot be used for data transmission. In burst-mode transmitters, it is important that the laser is capable of stabilizing quickly after being turned on, especially if the PHY layer requirements are tighter, as in the case of ITU-T G.984 GPON systems.

II.8 EPON Multi-Point Control Protocol

The Multi-Point Control Protocol (MPCP) was specified in [802.3ah] and extended to 10G-EPON in [802.3av] 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 the PON. 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 RTT (Round Trip Time), the allocated slots are always non-overlapping. This means that, upon their arrival at the OLT's 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:

- OLT MPCP sublayer is responsible for Discovery of the newly connected stations, their registration, measurement of the RTT, as well as scheduling and controlling the transmission from individual ONUs in the upstream channel;
- 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 a functionality through the respective GATE MPCPDU), as well as participation in the Discovery process.

The operation principle of the MPCP mechanism is relatively simple. The total available upstream channel bandwidth is divided into transmission units (typically termed slots) using the TDMA technique. Each slot 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 central OLT controller. 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 [802.3ah] and [802.3av]), 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 a MPCPDU. EPON specifications optionally allow the implementation 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 [802.3ah] and were 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 requires accounting for both the burst-mode delays and path delay variations between near and distant ONU's 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 protocol, 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 OLT controller to schedule at most 4 transmission slots at once (so called scheduling into the future), with the size of 216-1 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 reception of such an MPCPDU, the ONU updates its local clock index using the time-stamp field carried in the message body, effectively maintaining global synchronization with the OLT clock without the need for a separate clock signal. The scheduling information is parsed and processed accordingly. The result is 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, 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 inefficiency 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.

II.9 Unicast Traffic

Unicast traffic in EPONs is supported through emulation of P2P over P2MP medium by tagging the frames with a unicast LLID and a unicast MAC address. The tags allow filtering of the given frame at both the RS and MAC sublayers.

The objective of P2P Emulation (P2PE) mode is to achieve the same physical connectivity as in switched LAN, where all the stations are connected to a central switch using point-to-point links. In P2P emulation mode, the OLT must have N MAC ports (interfaces), one for each ONU and one more for the support of the broadcast frame delivery (the so called Single Copy Broadcast or SCB). During ONUs registration, a unique LLID value will be assigned to each ONU and each MAC port at the OLT will be assigned the same LLID as its corresponding ONU, providing a logical channel for the future communication between both end stations. In the downstream channel, a unicast frame broadcast to all ONUs over a PHY fiber plant is filtered at the RS sublayer by comparing the LLID included in the frame's preamble with the LLIDs supported by the given ONU. In the case of a match, the frame is announced to be addressed to the given ONU and passed to the MAC layer for further processing. Otherwise, it is discarded (not forwarded to the MAC sublayer). From the MAC sublayer perspective, it appears as if the frame was sent on a P2P link to only one of the ONUs.

In the upstream direction, the ONU will insert its assigned LLID in the preamble of each transmitted frame. P2PE in the OLT demultiplexes the frame to the proper MAC port based on the LLID.

II.10 IP Multicast

Like all Ethernet IEEE 802 protocols, EPON does not natively "understand" IP multicast. Delivery of the multicast services is accomplished by broadcasting to all ONUs where the Ethernet / IP level filtering can be carried out (depending on the implementation).

Multicast traffic can also be supported with the use of the Ethernet level filtering, where an ONU learns a set of "multicast" MAC addresses which are to be relayed to the UNI upon the reception of the data frames with such addresses on the PON interface. The process is typically heavily based on the IP level snooping (e.g., IGMP snooping), where the ONU identifies the MAC addresses of the multicast streams to be forwarded and includes this information in the Filtering Table (FT). Upon reception of frame with a broadcast LLID, MAC address-based filtering is carried out, and only the frames with the MAC addresses included in the FT will be relayed to the UNI. The remaining frames are filtered out, preventing their propagation in the customer network.

II.11 Broadcast Traffic

Broadcast traffic is inherently supported in the downstream channel of the EPON system, providing native support for any broadcast services. A broadcast data frame (regardless of whether it is a data frame or control frame) is designated always with a broadcast LLID address since it is transmitted by the SCB MAC instantiated in the OLT stack.

At the ONU, a broadcast frame will be filtered first at the RS sublayer (all frames with broadcast LLID are passed on to MAC layer for further processing) and only then processed at the MAC. A frame with a broadcast MAC address will be relayed to the customer UNI. A frame with the MAC Control address will be subject to (operator configurable) processing at the ONU.

II.12 P2P Emulation

EPON systems, due to the inherent requirements for P2PE over P2MP, require each downstream and upstream frame to be tagged with a network unique LLID on each PON. The number of LLIDs instantiated in a particular ONU has a significant impact on the system's performance and is one of the most vital design choices for a fully functional EPON system.

Appendix III Multi-Layer Switching

III.1 Multi-Layer Switching (MLS) and Classification

The DPoE specifications do not specify how to implement the interface between the EPON OLT and IP routing functions in the DPoE System. This section is informative only.

In a strictly layered model, IP access services might be constructed as a single [802.1] broadcast domain for each IP subnet. Since the DPoE architecture calls for [802.1ad] based architecture with EVCs for each service for each customer, such a basic model could create unnecessary complexity. Multi-Layer Switching (MLS) is a widely used technique in Ethernet switches and IP routers that combines Layer 2 (Ethernet) and Layer 3 (IP) forwarding into a single lookup. The MLS concept is informative for both DPoE System vendors and operators to understand the complexity of a multilayer forwarding model and how that model can be simplified.

III.2 MLS Forwarding

Although the service model uses full [802.1ad], it is possible to reduce the implementation complexity. In particular the DPoE System can use MLS to perform service-based forwarding. Such an implementation could, for example, build forwarding tables based on IP address, MAC address, S-VLAN ID, and C-VLAN ID (see Table 3 for example). Forwarding based on one or more of these parameters could eliminate the need to separately perform the [802.1ad] S component and C component bridge functions. Of course, each packet still needs to be fully encapsulated and forwarded.

III.3 IP(HSD) Implementation with MLS

MLS can be used to implement IP(HSD) forwarding over MEF implemented with [802.1ad] using a single step lookup. This can be accomplished by using an MLS implementation (as shown in Figure 9 to build a lookup and forwarding table as shown in Table 4 below.

Table 4 - IP(HSD) EVCs

Default Router (DR) Name	Default Router IP	Default Router MAC	ONU (eCM) MAC	Interface Type	eSAFE Type	eSAFE eDOCSIS IP	S-VLAN ID	C-VLAN ID
DR_1	1.1.1.254	<DR1 MAC>	<ONU A MAC>	CMCI	na	1.1.1.1	1	1
DR_1	1.1.1.254	<DR1 MAC>	<ONU A MAC>	LCI	eRouter	1.1.1.2	1	2
DR_2	2.2.2.254	<DR2 MAC>	<ONU A MAC>	LCI	sDVA	2.2.2.1	2	1
DR_2	2.2.2.254	<DR2 MAC>	<ONU A MAC>	LCI	eDVA	2.2.2.2	2	2
DR_1	1.1.1.254	<DR1 MAC>	<ONU B MAC>	CMCI	na	1.1.1.3	1	3
DR_2	2.2.2.254	<DR2 MAC>	<ONU B MAC>	LCI	eDVA	2.2.2.3	2	3

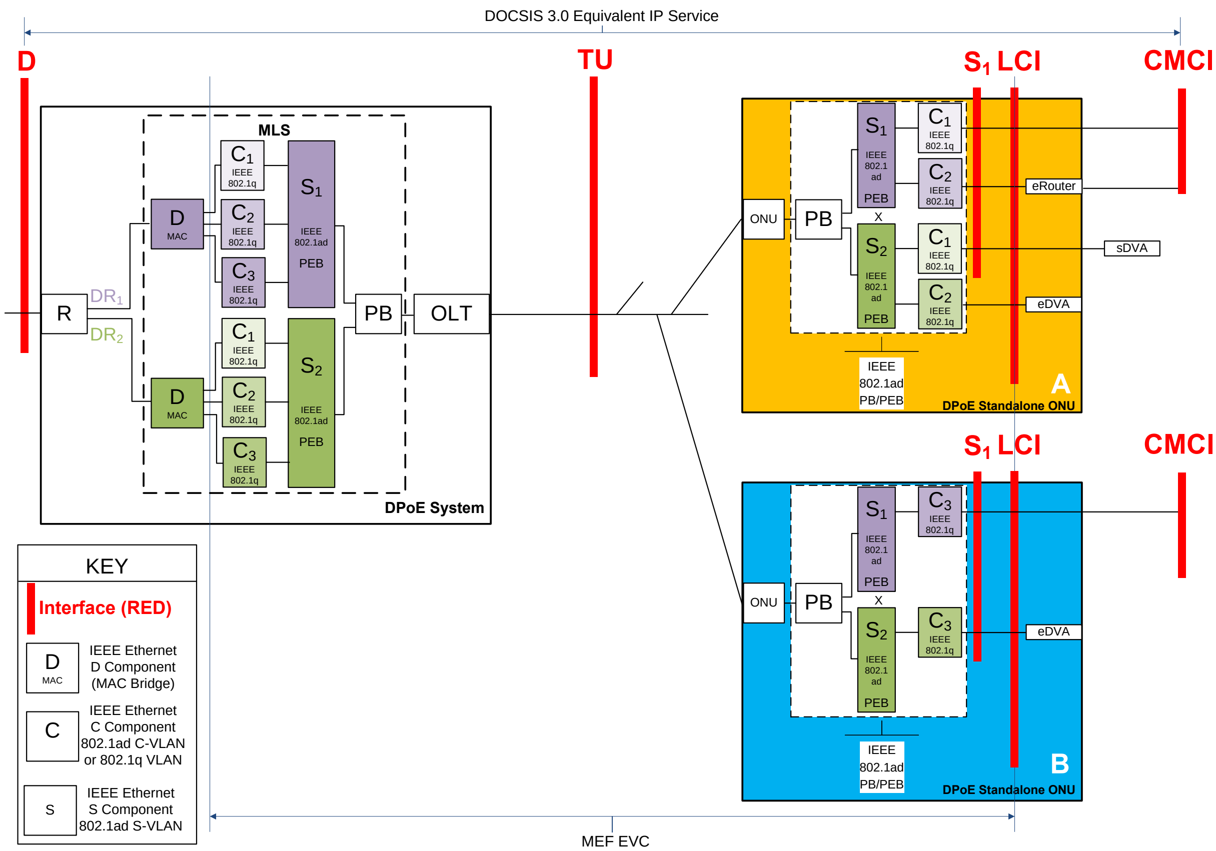

Figure 15 - IP(HSD) Implementation of MEF EVCs with MLS

III.4 IP Multicast

Eliminating the broadcast domain does not mean that broadcast technologies like Multicast cannot work. IP multicast can still work. This version of the DPoE specifications does not specify the implementation of IP multicast. IP multicast can be implemented in MLS by using any combination or permutation of S-VLAN ID, C-VLAN ID, source IP address, destination IP address, source MAC address, destination MAC address, or even other IP headers. As an example, each S-VLAN could be used as a "broadcast" domain for IP multicast. This could be accomplished by configuring MLS to forward all traffic from a source multicast address to all C-VLAN IDs on a given S-VLAN for which IGMP joins are received (for that source address with PIM-SSM). A single multicast proxy running on the IP router can perform all of the functionality required. Instead of transmitting the traffic on each EVC, the traffic could be transmitted on the S-VLAN. With appropriate classifiers on the DPoE System and DPoE ONUs, such a system could effectively deliver IP multicast using EPON broadcast capabilities while still operating only two port bridge groups. In this particular example, the S-VLAN provides the multiport bridge group that can provide broadcast, but only for IP multicast and only when joins are properly proxied (received and responded to).

III.5 MAC Optional

Once an EVC is established, the MAC of the eSAFE or eDOCSIS device (on an LCI) or the MAC of the DOCSIS CPE (on a CMCI) is visible to the DPoE System where the S-VLAN and C-VLAN terminate before forwarding on the IP router. In other words, the default router (DR) interface on the Router (R) are in the same [802.1d] bridge group (called a broadcast domain) and IP and MAC associative protocols like ARP, Reverse-ARP (RARP), and Neighbor Discovery Protocol (NDP) can run utilizing the Layer 2 ([802.1d]) adjacency to associate IP and MAC addresses. If MLS is operated, the forwarding engine on a DPoE System does not require ARP and could function by passively learning the (eDOCSIS or CPE) device address with DHCP. In short, MLS would allow for what is called Layer 3 switching. That is forwarding to Layer 2 connected devices based on a forwarding table that associates IP addresses directly with (logical or physical ports). In a DPoE System implementation, an MLS implementation could forward IP traffic to an S-VLAN ID + C-VLAN ID pair without respect to MAC addresses. Although this technology is not currently used in DOCSIS specifications (nor is it required in the DPoE specifications), it is common among Layer 3 switches, which are indistinguishable from traditional double lookup routers. IP forwarding without MAC addresses is also common on other circuit-based technologies such as Packet Over SONET, PPP, and SLIP. In this context, the EVC is a point-to-point circuit. Because there is one and only one destination on the opposite side of any EVC, IP forwarding to the next hop can be performed without a Layer 2 (MAC) address lookup. Another way of saying this is that the use of a MAC is really for forwarding on interfaces where there are multiple receivers (devices). By using a point-to-point (or circuit like) model such as the EVC, MAC addresses are not necessary.

Appendix IV Example LLID Requirements

The minimum number of LLIDs (bi-directional unicast LLIDs) required by the DPoE specifications can be calculated based on the number of S interfaces on the DPoE ONU. The S interface is described in the specifications above. As specified, the formula for the calculation is:

Minimum number of LLIDs a DPoE ONU MUST have = $\min_8 [\![(S+1)]\!]$

Number of LLIDs a DPoE ONU SHOULD have = $\min_8 [\![((S \text{ x } 2)+1)]\!]$

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8[(S\text{-}LCI)x\ 2+(LCI\ x\ 2)+1]$

Assume each eDVA is 2 POTS or T.38 (LCI sub-ports). Assume each eRouter has a single Ethernet (LCI sub-port).

Table 5 - A summary of possible DPoE ONU LLID Requirements

	Ethernet Ports	LCI Ports	LCI Sub-Ports	Total Ports	LLID Calculated	LLID MUST	LLID SHOULD	LLID MAY
			(LCI x 2)	(S)	S + 1	[S + 1] or 8 min.	S x 2 + 1 or 8 min.	(S – LCI) x2 + (LCI x 2) + 1 or 8 min.
Bridge ONU in 6 port switch	5			5	6	8	8	11
Bridge ONU in 24 port switch	23			23	24	8	49	51
1 Port DPoE ONU	1	1	0	1	2	8	8	8
2 Port DPoE ONU	2	2	0	2	3	8	8	8
1 Port DPoE ONU w/ eRouter	1	1	1	2	3	8	8	8
2 Port DPoE ONU w/ eRouter	2	2	1	3	3	8	8	8
1 Port DPoE ONU w/eDVA	1	1	2	3	3	8	8	8
2 Port DPoE ONU w/eDVA	2	1	2	3	4	8	8	9
1 Port DPoE ONU w/ eRouter and eDVA	1	2	3	2	4	8	8	8
2 Port DPoE ONU w/ eRouter and eDVA	2	2	2	3	5	8	8	10
4 Port DPoE ONU	4	0	0	4	5	8	9	9

	Ethernet Ports	LCI Ports	LCI Sub-Ports	Total Ports	LLID Calculated	LLID MUST	LLID SHOULD	LLID MAY
			(LCI x 2)	(S)	S + 1	[S + 1] or 8 min.	S x 2 + 1 or 8 min.	(S – LCI) x2 + (LCI x 2) + 1 or 8 min.
4 Port DPoE ONU w/ eRouter and eDVA	4	2	2	6	7	8	12	15
4 port DPoE ONU w/ 16 line eDVA	4	1	16	5	6	8	11	41
7 Port DPoE ONU	7	0	0	7	8	8	15	15
16 Port DPoE ONU	16	0	0	16	17	17	17	17

IV.1 Example 1. DPoE ONU #1

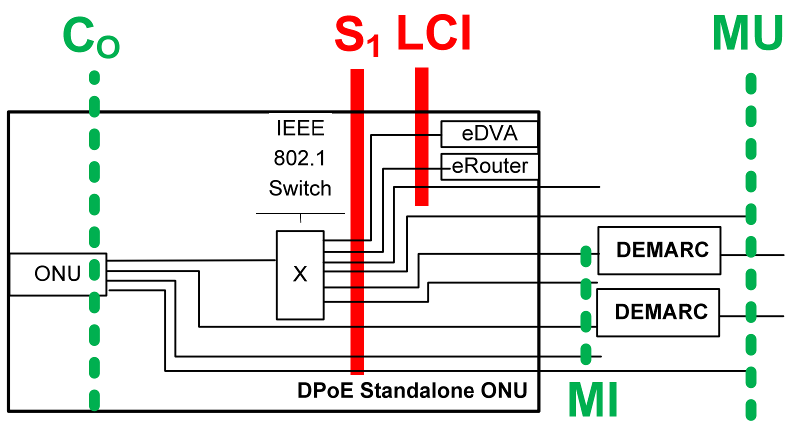

Figure 16 - Example 1. DPoE ONU #1

For the example shown in Figure 16, there are two S interfaces serving eDOCSIS embedded subsystems, one dedicated S_1 port for an LCI interface, and six (6) additional Ethernet interfaces operable as MU or MI interfaces. Assume that the eDVA has 2 ports and the eRouter has 2 ports.

S1 Count = LCI + MI + MU
S1 Count = 3 + 4 + 2 = 9

MUST

Minimum number of LLIDs a DPoE ONU MUST have = $\min_8[(S+1)]$

Minimum number of LLIDs a DPoE ONU MUST have = 9 + 1

Minimum number of LLIDs a DPoE ONU MUST have = 10

SHOULD

Number of LLIDs a DPoE ONU SHOULD have = min8[(S x 2)+ 1)]

Number of LLIDs a DPoE ONU SHOULD have = min8[((9 x 2)+ 1)]

Number of LLIDs a DPoE ONU SHOULD have = min8[(19)]

Number of LLIDs a DPoE ONU SHOULD have = 19

MAY

LCI Sub-interfaces = 2 + 2 = 4 (eDVA = 2; eRouter = 2)

LCI direct = 1

MI = 4

MU = 2

Total = 11

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8$[*(S-LCI)*x 2+*(LCI x 2)*+ 1]

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8$[(9-3)x 2+(3 x 2)+ 1]

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8$[12+6+1]

Number of LLIDs a DPoE ONU MAY ideally have = 19

IV.2 Example 2. DPoE ONU #2

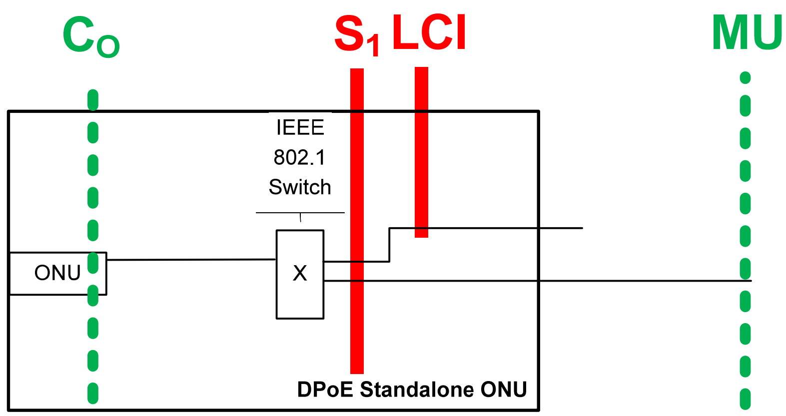

Figure 17 - Example 2. DPoE ONU #2

For the example shown in Figure 17, there are two S interfaces. One is an LCI interface and one S interface is connected to an Ethernet physical port, operable as MI or MU interfaces (depicted as an MU interface).

MUST

Minimum number of LLIDs a DPoE ONU MUST have = $\min_8$[(*S*+1)]

Minimum number of LLIDs a DPoE ONU MUST have = $\min_8$[(2+1)]

Minimum number of LLIDs a DPoE ONU MUST have = $\min_8$[(3)]

Minimum number of LLIDs a DPoE ONU MUST have = 8

The specifications require the GREATER of the computed minimum or eight. Since the computed minimum is less than 8, the minimum MUST be 8).

SHOULD

Number of LLIDs a DPoE ONU SHOULD have =$\min_8$[((*S* x 2)+ 1)]

Number of LLIDs a DPoE ONU SHOULD have =$\min_8$[((3 x 2)+ 1)]

Number of LLIDs a DPoE ONU SHOULD have =$\min_8$[(7)]

Number of LLIDs a DPoE ONU SHOULD have =8

MAY

LCI Sub-interfaces = 1 (LCI direct = 1)

MI or MU = 1

Total = 2

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8$[*(S-LCI)x 2+(LCI x 2)*+ 1]

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8$[(2-1)x 2+(1 x 2)+ 1]

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8$[4+2+1]

Number of LLIDs a DPoE ONU MAY ideally have = 8

IV.3 Example 3. DPoE ONU #3

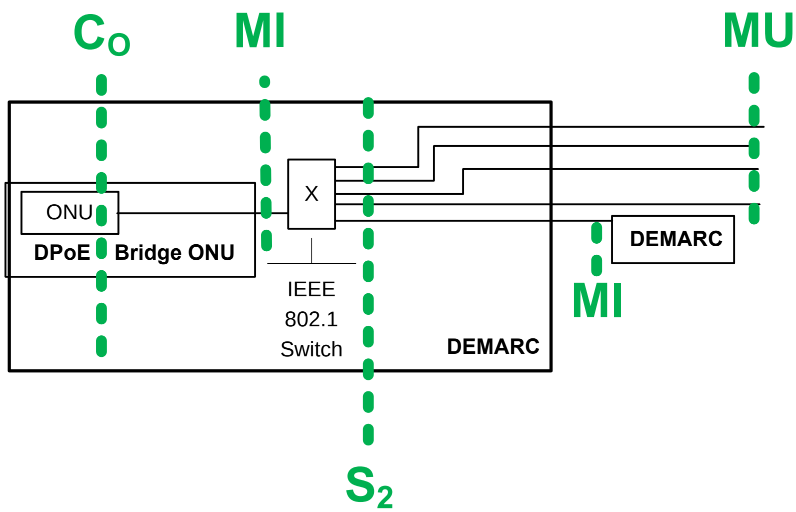

Figure 18 - Example 3. DPoE ONU #3

For the example shown in Figure 18, there are five (5) S interfaces connected to an Ethernet physical ports on the DEMARC device. Each interface is operable as an MI or an MU interface.

MUST

Minimum number of LLIDs a DPoE ONU MUST have = $\min_8$[(*S*+1)]

Minimum number of LLIDs a DPoE ONU MUST have = $\min_8$[(5+1)]

Minimum number of LLIDs a DPoE ONU MUST have = $\min_8$[(6)]

Minimum number of LLIDs a DPoE ONU MUST have = 8

The specifications require the GREATER of the computed minimum or eight. Since the computed minimum is less than 8, the minimum MUST be 8.

SHOULD

Number of LLIDs a DPoE ONU SHOULD have = 5 x 2 + 1

Number of LLIDs a DPoE ONU SHOULD have = 10 + 1

Number of LLIDs a DPoE ONU SHOULD have = 11

MAY

MI or MU = 5

Total = 5

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8$[*(S-LCI)x 2+(LCI x 2)*+ 1]

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8$[(5-0)x 2+(0 x 2)+ 1]

Number of LLIDs a DPoE ONU MAY ideally have = $\min_8$[10+0+ 1]

Number of LLIDs a DPoE ONU MAY ideally have = 11

When there are no LCI interfaces, the "SHOULD" and "MAY" specifications will always be the same.

Appendix V 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 Meiyan	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