

DOCSIS® Provisioning of EPON Specifications

DPoEv2.0

DPoE Architecture Specification

DPoE-SP-ARCHv2.0-I08-230322

ISSUED

Notice

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

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

© Cable Television Laboratories, Inc., 2011 – 2023

DISCLAIMER

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

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

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

Document Status Sheet

Document Control Number:	DPoE-SP-ARCHv2.0-I08-230322			
Document Title:	DPoE Architecture Specification			
Revision History:	I01 – Released 10/04/12 I02 – Released 08/07/14 I03 – Released 03/19/15 I04 – Released 06/11/15 I05 – Released 06/02/16 I06 – Released 02/28/18 I07 – Released 02/13/19 I08 – Released 03/22/23			
Date:	March 22, 2023			
Status:	Work in Progress	Draft	Issued	Closed
Distribution Restrictions:	Author Only	GL/Member	GL/Member/Vendor	Public

Key to Document Status Codes

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

Trademarks

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

Contents

1	INTRODUCTION	7
1.1	DPoE Technology Introduction	7
1.2	Scope	8
1.3	DPoE Architecture Specification Goals.....	8
1.4	Requirements	8
1.5	DPoE Version 2.0 Specifications.....	9
1.6	Reference Architecture	10
1.7	DPoE Interfaces and Reference Points	10
2	REFERENCES	14
2.1	Normative References.....	14
2.2	Informative References.....	15
2.3	Reference Acquisition.....	17
3	TERMS AND DEFINITIONS	18
3.1	DPoE Network Elements.....	18
3.2	Other Terms	20
4	ABBREVIATIONS AND ACRONYMS.....	21
5	DPOE SERVICE REQUIREMENTS.....	25
5.1	Business Services Overview	25
5.2	Service Requirements	26
5.2.1	<i>Introduction</i>	26
5.2.2	<i>Private Data Services</i>	26
5.2.3	<i>IP (Public IP/Internet)</i>	26
5.2.4	<i>Voice Services</i>	26
5.2.5	<i>Vertical Markets</i>	26
5.3	Single Tenant Businesses.....	27
5.4	Multiple Tenant Businesses	27
5.5	DPoE Fully Automated and Partially Automated Services	27
5.5.1	<i>Private Data Services</i>	27
5.5.2	<i>Internet Access Service</i>	28
5.5.3	<i>Internet Transit Service</i>	28
5.5.4	<i>Internet Peering Function</i>	28
5.6	Recommendations for DPoE as a Metro Ethernet or IP Transport.....	28
5.6.1	<i>Voice Services</i>	28
5.6.2	<i>Wholesale Ethernet</i>	29
5.6.3	<i>T1 or T3 Circuit Emulation</i>	29
5.6.4	<i>IP-PBX</i>	29
5.6.5	<i>Cell Tower Backhaul</i>	29
5.6.6	<i>Video Distribution and Transport</i>	30
5.6.7	<i>Storage Network Services</i>	30
5.6.8	<i>Cloud Services</i>	30
5.6.9	<i>Vertical Markets</i>	30
6	ARCHITECTURAL FOUNDATION.....	32
6.1	Forwarding.....	32
6.1.1	<i>Ethernet Virtual Connections</i>	32
6.1.2	<i>Metro Ethernet Services</i>	32
6.1.3	<i>Internet Services</i>	32
6.2	Service Integration.....	33
6.3	Multipoint Architecture	33

6.4	DOCSIS Emulation	33
6.5	eDOCSIS	34
6.6	EPON OAM	35
7	REQUIREMENTS	36
7.1	Introduction	36
7.2	Architectural Elements	36
7.2.1	<i>DPoE Network Elements</i>	36
7.2.2	<i>DPoE Network Definition</i>	36
7.2.3	<i>DPoE System Definition</i>	36
7.2.4	<i>DPoE ONU Definition</i>	36
7.2.5	<i>DPoE Standalone ONU Definition</i>	37
7.2.6	<i>DPoE Bridge ONU Definition</i>	37
7.3	Interface and Reference Point Detailed Descriptions	37
7.3.1	<i>DPoE Interfaces</i>	37
7.3.2	<i>DPoE Reference Interfaces and Reference Points</i>	39
7.4	Architectural Requirements	40
7.4.1	<i>Ethernet Service Based Architecture</i>	40
7.5	Service Model	41
7.5.1	<i>Service Multiplexing Model</i>	41
7.5.2	<i>Classification Reference Points</i>	41
7.6	IP(HSD) Forwarding Model	42
7.7	L2HSD Forwarding Model	42
7.8	eDOCSIS	42
7.8.1	<i>eDOCSIS on S-ONU</i>	42
7.8.2	<i>eSAFE Discovery</i>	42
7.8.3	<i>vCM and eCM</i>	43
7.8.4	<i>eDOCSIS Applicability</i>	43
7.9	Power Saving	43
7.10	Optical Link Protection	44
8	EVPL, E-LAN, AND E-TREE SERVICES	45
8.1	E-Line Service	46
8.2	E-LAN and E-Tree Services	46
8.3	Supported Forwarding Models	46
8.3.1	<i>Fully Centralized</i>	46
8.3.2	<i>Fully Distributed</i>	46
8.3.3	<i>Hybrid Centralized-Distributed</i>	47
8.4	EVPL, E-LAN, E-Tree Service Provisioning	47
8.4.1	<i>DOCSIS-L2VPN Service Provisioning Model</i>	47
8.4.2	<i>[RFC 6074]-Based Service Provisioning</i>	49
8.4.3	<i>Associating Service Delimiting Tags to an MPLS Service</i>	51
8.4.4	<i>Simplified MPLS Provisioning Mode for Encapsulation Mode</i>	52
8.5	Local VSI and PW Forwarder Configuration Management Requirements	53
8.5.1	<i>Adding Attachment Circuits to an Active VSI or PW Instance on a DPoE System</i>	53
8.5.2	<i>Modifying an Active VSI or PW Instance on a DPoE System</i>	54
8.5.3	<i>Removing Attachment Circuits from an Active VSI or PW Instance on a DPoE System</i>	54
8.6	L2VPN Layer 2 Forwarding Services	54
8.7	IP Transport Introduction	55
8.7.1	<i>General VPWS/VPLS Forwarding</i>	55
8.7.2	<i>E-Tree Forwarding</i>	56
APPENDIX I ACKNOWLEDGEMENTS		61
APPENDIX II REVISION HISTORY		62
II.1	Engineering Changes for DPoE-SP-ARCHv2.0-I02-140807	62

II.2	Engineering Change for DPoE-SP-ARCHv2.0-I03-150319.....	62
II.3	Engineering Changes for DPoE-SP-ARCHv2.0-I04-150611	62
II.4	Engineering Changes for DPoE-SP-ARCHv2.0-I05-160602	62
II.5	Engineering Change for DPoE-SP-ARCHv2.0-I06-180228.....	62
II.6	Engineering Change for DPoE-SP-ARCHv2.0-I07-190213.....	62
II.7	Engineering Change for DPoE-SP-ARCHv2.0-I08-230322.....	63

Figures

Figure 1 - DPoEv2.0 Reference Architecture.....	10
Figure 2 - DPoEv2.0 Interfaces and Reference Points.....	11
Figure 3 - D-ONU Types.....	19
Figure 4 - DPoE Network Elements	19
Figure 5 - eDOCSIS Reference Model (from Figure 5-1 of [eDOCSIS])	34
Figure 6 - eCM Functionality in an eDOCSIS Device Split Between DPoE System and S-ONU	35
Figure 7 - S-ONU that is an eDOCSIS Device, with vCM in DPoE System	35
Figure 8 - DPoE Specifications Support Both Embedded SAFEs and External CPE	43
Figure 9 - ELAN-MPLS DPoE Model	45
Figure 10 - Encapsulation NSI Extension.....	47
Figure 11 - L2VPN Example Config.....	48
Figure 12 - Example CM Configuration File.....	49
Figure 13 - VSI New TLV Object Model.....	50
Figure 14 - VE Reference Model.....	55
Figure 15 - E-Tree Forwarding.....	56
Figure 16 - DPoE-PBB-DEMARC.....	57
Figure 17 - PBB-D-ONU.....	58
Figure 18 - PBB Upstream Forwarding.....	59
Figure 19 - PBB Downstream Forwarding.....	60

Tables

Table 1 - DPoEv2.0 Series of Specifications.....	9
Table 2 - DPoEv2.0 Interface and Reference Point Descriptions.....	12
Table 3 - DPoE Network Element Abbreviations.....	36
Table 4 - Examples of BP-ONUs	37
Table 5 - L2 Forwarding Modes.....	54

1 INTRODUCTION

DOCSIS Provisioning of EPON (DPoE) version 2.0 specifications are a joint effort of Cable Television Laboratories (CableLabs), cable operators, vendors, and suppliers to support EPON technology using existing DOCSIS-based back office systems and processes. DPoEv2.0 specifications augment the DPoE v1.0 specifications to provide requirements for additional service capabilities and corresponding provisioning and network management capabilities.

Ethernet PON (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.

DPoE specifications are focused on DOCSIS-based provisioning and operations of Internet Protocol (IP) using DOCSIS Internet service (which is typically referred to as High Speed Data (HSD)), or IP(HSD) for short, and Metro Ethernet services as described by Metro Ethernet Forum (MEF) standards. 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) together 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. 2G-EPON, as described Annex A of [DPoE-PHYv2.0], uses the same 1G upstream as 1G-EPON (operates at the effective rate of 1 Gbps), but provides a 2G downstream (operates at the effective rate of 2 Gbps). With the exception of requirements specified in Annex A of [DPoE-PHYv2.0], 2G-EPON is expected to meet all of the requirements specified for 1G-EPON. 10G-EPON technology is available and is backwards compatible with 1G-EPON. A 1G-EPON network can be incrementally upgraded to 10G-EPON, adding or replacing ONUs as business needs require. 1G-EPON, 2G-EPON, and 10G-EPON are compatible with [SCTE 174].

1G-EPON, originally defined in [802.3ah], 2G-EPON part of the CCSA specifications ([YD/T 1993.1-2009], and [YD/T 1993.2-2010]), and 10G-EPON [802.3av] respectively, support a point-to-multipoint architecture with a centralized controller called an Optical Line Terminal (OLT) and distributed low cost Layer 2 ONUs. The basic service mapping architecture in EPON is to map Ethernet (or IP) frame header information (e.g., addresses, IP Differentiated Service Code Points, Ethernet Q tag, S-VLAN/C-VLAN ID, ISID, bridge address, etc.) to a logical circuit called a Logical Link Identifier (LLID) in [802.3ah]. The service mapping function in DPoE specifications is similar to that used in DOCSIS specifications. Both DOCSIS and DPoE networks rely on a centralized scheduler though EPON utilizes an LLID which functions like a SID in DOCSIS to support unicast, broadcast, and multicast.

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

1.2 Scope

This specification describes the version 2.0 architecture for DPoE Networks.

Existing business services include one or more DOCSIS IP services, baseband and broadband Ethernet services over coaxial cable, IP and Ethernet over fiber with baseband and broadband (Course Wavelength Division Multiplexing (CWDM) and Dense Wavelength Division Multiplexing (DWDM)), Broadband Passive Optical Network (BPON), Ethernet Passive Optical Network (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 processes 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 E-Line, E-LAN, and E-Tree services
- IP(HSD) Internet Service as defined by DOCSIS 3.0.

An operator can offer IP-based services over the Metro Ethernet service platform by treating IP as an application over the DPoE Network.

Other services can be operated "over-the-top" of the above services, but the provisioning, operations, and other requirements for such services are not specified in this version of DPoE specifications.

1.3 DPoE Architecture Specification Goals

The DPoE Architecture specification accomplishes the following objectives:

- Identify and document the common requirements for services for business customers over EPON.
- Define and describe the interfaces and reference points as part of the DPoE reference architecture.
- Define and describe the elements that collectively form a DPoE Network.
- Communicate the architectural foundation on which other DPoE specifications depend.
- Define requirements for services supported using DPoEv2.0 specifications.

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 DPoE Version 2.0 Specifications

A list of the specifications included in the DPoEv2.0 series is provided in Table 1. For further information, please refer to <http://www.cablelabs.com/specs/specification-search/?cat=dpoe&scat=dpoe-2-0>.

Table 1 - DPoEv2.0 Series of Specifications

Designation	Title
DPoE-SP-ARCHv2.0	DPoE Architecture Specification
DPoE-SP-OAMv2.0	DPoE OAM Extensions Specification
DPoE-SP-PHYv2.0	DPoE Physical Layer Specification
DPoE-SP-SECv2.0	DPoE Security and Certificate Specification
DPoE-SP-IPNEv2.0	DPoE IP Network Element Requirements
DPoE-SP-MULPIv2.0	DPoE MAC and Upper Layer Protocols Interface Specification
DPoE-SP-MEFv2.0	DPoE Metro Ethernet Forum Specification
DPoE-SP-OSSiv2.0	DPoE Operations and Support System Interface Specification

1.6 Reference Architecture

The DPoE reference architecture shown in Figure 1 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 headend or hub site, and the DPoE ONU (D-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 the server elements have connectivity through an IP (TCP/IP) network. Transport of bearer traffic, and (in some cases) Layer 2 OAM Protocol Data Units (PDUs) are available through either IP or Layer 2 Ethernet-based Network Interfaces.




Figure 1 - DPoEv2.0 Reference Architecture

1.7 DPoE Interfaces and Reference Points

The DPoE interfaces and reference points shown in Figure 2 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 subsystems. The reference points have protocols that run across them, or have a common format of bearer traffic (with no signaling protocol). All 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 unidirectional for upstream (C₀) or downstream (C_s) classification, respectively.




Figure 2 - DPoEv2.0 Interfaces and Reference Points

Table 2 - DPoEv2.0 Interface and Reference Point Descriptions

Interface or Reference Point		Interface or Reference Point Description
MN		MN is a logical concept used for the specification of requirements for MEF INNI that apply to both MN _E and MN _I . MN logically provides the equivalent function of a MEF INNI or L2VPN NSI. It is an NNI for Metro Ethernet services only.
	MN _E	The MN _E (MEF INNI External) interface is a substitute for the MN reference interface from DPoE version 1.0 specifications. 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.
	MN _I	The MN _I reference interface is used to describe the virtual interface between an OLT and a VPLS Virtual Switch Instance (VSI). In particular, it is used to describe the requirements for stitching VSIs to DPoE System and OLT [802.1] components such as [802.1Q] bridge groups, [802.1ad] S-VLAN or C-VLAN (S-component or C-component), or [802.1ad] I-BEB (I-component) or B-BEB (B-component) backbone edge bridges. The DPoE System stitches VPLS and VPWS transport and forwarding for Metro Ethernet Services between the D interface and the MN _I reference interface ¹ .
D		The D interface is the DOCSIS IP NNI interface. It is an operator network-facing interface, sometimes called a 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, IP/MPLS/VPLS traffic, and IP/MPLS/VPWS traffic.
TU		The TU interface is the interface between the DPoE System and the D-ONU.
TUL		The TUL interface is a virtual interface representing a logical EPON on an ODN. Each ODN has at least one TUL, and each TUL represents a MAC domain.
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 D-ONU upstream classifier.
	C _S	The C _S reference point is used for explanation of traffic ingress to a DPoE System downstream classifier.
S		The S interface is an IEEE 802 interface. The S interface may be an internal interface, (such as an [802.3] GMII or XGMII interface implemented in the form of a SERDES) or an external interface exposed in a BB-ONU or S-ONU for connection to other devices. The S interface may be expressed in the form of CMCI, LCI, MU, MI, or other forms of interfaces internal to the ONU.
LCI		The Logical CPE Interface (LCI) interface is an eDOCSIS interface as defined in [eDOCSIS]. eSAFEs are connected to LCI interfaces.
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. Customer Premise Equipment (CPE) is connected to CMCI interfaces.

¹ MN_I is required for IP-based forwarding and transport of Metro Ethernet services with DPoE in order to provide MEF E-LAN and E-TREE services described in DPoE version 2.0. While these services can be constructed with MN_E, these specifications do not describe the process to do so.

Interface or Reference Point	Interface or Reference Point Description
MI	<p>MI is an S interface that operates as a MEF INNI with additional requirements as specified in [DPoE-MEFv2.0]. The MI interface is an [802.3] interface (or reference point) between a D-ONU and a DEMARC device.</p> <ul style="list-style-type: none"> • A D-ONU that provides a MEF INNI has an MI interface. • A D-ONU can have MU as an interface and an MI reference point on different S interfaces in a single D-ONU. <p>DEMARC devices are connected to MI interfaces.</p>
MU	<p>MU is an S interface (or S reference interface) that operates as a MEF UNI. The MU reference interface is an [802.3] interface (or reference point) between a D-ONU or a DEMARC device and a customer's equipment.</p> <ul style="list-style-type: none"> • A D-ONU that directly provides a MEF UNI (MU) interface has MU as an interface. • A D-ONU can have MU as an interface and an MI reference point on different S interfaces in a single D-ONU. <p>Customer Edge (CE) devices are connected to MU interfaces.</p>

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.

In this specification, terms "802.1ad" and "802.1ah" are used to indicate compliance with the [802.1ad] and [802.1ah] standards, respectively, now incorporated as part of [802.1Q]. For all intents and purposes, claiming compliance to [802.1Q], [802.1ad] or [802.1ah] in the scope of this specification will be treated as claiming compliance to IEEE Std 802.1Q-2011. Unless otherwise stated, claiming compliance to 802.1Q-2005 requires a specific date reference.

- [1904.1A] IEEE Std 1904.1-2016, IEEE Standard for Service Interoperability in Ethernet Passive Optical Networks (SIEPON), Package A, draft D2.0.
- [802.1] Refers to entire suite of IEEE 802.1 standards unless otherwise specified.
- [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. Former amendment to 802.1Q, now part of 802.1Q-2011.
- [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. Former amendment to 802.1Q, now part of 802.1Q-2011.
- [802.1Q] IEEE Std 802.1Q-2018, IEEE Standard for Local and Metropolitan Area Networks – Bridges and Bridged Networks, July 2018.
- [802.3] IEEE Std 802.3-2012, IEEE Standard for Ethernet, December 2012.
- [802.3ah] IEEE Std 802.3ah-2004, IEEE Standard for Information Technology – Telecommunications and Information Systems – Local and Metropolitan Area Networks – Specific Requirements, Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications, Amendment: Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks, now part of [802.3].
- [802.3av] IEEE Std 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, now part of [802.3].
- [CMCIv3.0] Data-Over-Cable Service Interface Specifications, Cable Modem to Customer Premise Equipment Interface Specification, CM-SP-CMCIv3.0-I03-170510, March 10, 2017, Cable Television Laboratories, Inc.
- [DPoE-IPNEv2.0] DOCSIS Provisioning of EPON, IP Network Element Requirements, DPoE-SP-IPNEv2.0-I08-230322, March 22, 2023, Cable Television Laboratories, Inc.
- [DPoE-MEFv2.0] DOCSIS Provisioning of EPON, Metro Ethernet Forum Specification, DPoE-SP-MEFv2.0-I07-230322, March 22, 2023, Cable Television Laboratories, Inc.
- [DPoE-MULPIv2.0] DOCSIS Provisioning of EPON, MAC and Upper Layer Protocols Interface Specification, DPoE-SP-MULPIv2.0-I14-230322, March 22, 2023, Cable Television Laboratories, Inc.

- [DPoE-OAMv2.0] DOCSIS Provisioning of EPON, OAM Extensions Specification, DPoE-SP-OAMv2.0-I15-230322, March 22, 2023, Cable Television Laboratories, Inc.
- [DPoE-OSSIV2.0] DOCSIS Provisioning of EPON, Operations and Support System Interface Specification, DPoE-SP-OSSIV2.0-I13-230322, March 22, 2023, Cable Television Laboratories, Inc.
- [DPoE-PHYv2.0] DOCSIS Provisioning of EPON, Physical Layer Specification, DPoE-SP-PHYv2.0-I07-230322, March 22, 2023, Cable Television Laboratories, Inc.
- [DPoE-SECv2.0] DOCSIS Provisioning of EPON, Security and Certificate Specification, DPoE-SP-SECv2.0-I07-230322, March 22, 2023, Cable Television Laboratories, Inc.
- [eDOCSIS] Data-Over-Cable Service Interface Specifications, eDOCSIS Specification, CM-SP-eDOCSIS-I31-220831, August 31, 2022, Cable Television Laboratories, Inc.

2.2 Informative References

This specification uses the following informative references.

- [802.1ag] IEEE Std 802.1ag-2007, IEEE Standard for Local and Metropolitan Area Networks – Virtual Bridged Local Area Networks Amendment 5: Connectivity Fault Management, December 2007.
- [802.1ax] IEEE Std 802.1ax-2008, IEEE Standard for Local and Metropolitan Area Networks – Link Aggregation, January 2008.
- [802.3as] IEEE Std 802.3asTM2006, 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, now part of [802.3].
- [DOCSIS] Refers to DOCSIS 3.0 unless otherwise specified.
- [eRouter] Data-Over-Cable Service Interface Specifications, eRouter Specification, CM-SP-eRouter-I21-220209, February 9, 2022, Cable Television Laboratories, Inc.
- [G.805] ITU-T Recommendation G.805 (03/2000), Generic Functional Architecture of Transport Networks.
- [L2VPN] Data-Over-Cable Service Interface Specifications, Layer 2 Virtual Private Networks, CM-SP-L2VPN-I16-220328, March 28, 2022, Cable Television Laboratories, Inc.
- [MEF 6] Metro Ethernet Forum, MEF 6.1 Ethernet Services Definitions, Phase 2, April 2008.
- [MEF 9] Metro Ethernet Forum, Abstract Test Suite for Ethernet Services at the UNI, October 2004.
- [MEF 14] Metro Ethernet Forum, Abstract Test Suite for Traffic Management Phase 1, November 2005.
- [MEF 21] Metro Ethernet Forum, Service OAM and Requirements Framework, Phase 1, April 2007.
- [MEF 26] Metro Ethernet Forum, External Network to Network Interface (ENNI) – Phase 1, January 2010.
- [MULPIv3.0] Data-Over-Cable Service Interface Specifications, MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.0-C01-171207, December 7, 2017, Cable Television Laboratories, Inc.
- [OSSIV3.0] Data-Over-Cable Service Interface Specifications, Operations Support System Interface Specification, CM-SP-OSSIV3.0-C01-171207, December 7, 2017, Cable Television Laboratories, Inc.
- [PHYv3.0] Data-Over-Cable Service Interface Specifications, Physical Layer Specification, CM-SP-PHYv3.0-C01-171207, December 7, 2017, Cable Television Laboratories, Inc.

- [RFC 1918] IETF RFC 1918, Address Allocation for Private Internets, February 1996.
- [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, June 2000.
- [RFC 3418] IETF RFC 3418/STD0062, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP), June 2000.
- [RFC 4188] IETF RFC 4188, Definitions of Managed Objects for Bridges, September 2005.
- [RFC 4293] IETF RFC 4293, Management Information Base for the Internet Protocol (IP), April 2006.
- [RFC 4364] IETF RFC 4364, BGP/MPLS IP Virtual Private Networks (VPNs), February 2006.
- [RFC 4448] IETF RFC 4448, Encapsulation Methods for Transport of Ethernet over MPLS Networks, April 2006.
- [RFC 4761] IETF RFC 4761, Virtual Private LAN Service (VPLS), Using BGP for Auto-Discovery and Signaling, January 2007.
- [RFC 4762] IETF RFC 4762, Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling, January 2007.
- [RFC 5129] IETF RFC 5129, Explicit Congestion Marking in MPLS, January 2008.
- [RFC 6074] IETF RFC 6074, Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs), January 2011.
- [SCTE 174] ANSI/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-C01-171207, December 7, 2017, 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 May 12, 2001.
- [YD/T 1993.1-2009] CCSA Technical Requirements for Access Network 2Gbps Ethernet Passive Optical Network (2G-EPON) Part1: Compatible Mode
- [YD/T 1993.2-2010] CCSA Technical Requirements for Access Network 2Gbps Ethernet Passive Optical Network (2G-EPON) Part 2: Co-exist Mode

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>
- CCSA, China Communications Standards Association, 52# Hua Yuan Bei Road, Haidian District, Beijing, P.R.China 100083, Telephone+86-10-62302645, Fax+86-10-62301849, Internet: <http://www.ccsa.org.cn/english/>
- 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>
- ITU: International Telecommunications Union (ITU), <http://www.itu.int/home/contact/index.html>
- MEF: Metro Ethernet Forum, 6033 W. Century Blvd, Suite 830, Los Angeles, CA 90045
Phone +1-310-642-2800; Fax +1-310-642-2808. Internet: <http://metroethernetforum.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 Network Elements

DPoE Network	This term means all the elements of a DPoE implementation, including at least one DPoE System and one or more D-ONUs connected to that DPoE System.
DPoE System	This term refers to the set of subsystems within the hub site that provides the functions necessary to meet DPoE specification requirements.
DPoE ONU (D-ONU)	This term means a DPoE-capable ONU that complies with all the DPoE specifications. There are two logical types of D-ONUs. These are the DPoE Standalone ONU (S-ONU) and the DPoE Bridge ONU (B-ONU). Requirements specified for a D-ONU must be met by all ONUs.
DPoE Standalone ONU (S-ONU)	This term means a D-ONU that provides all the functions of a B-ONU and also provides at least one CMCI port. An S-ONU can optionally have one or more eSAFEs.
DPoE Bridge ONU (B-ONU)	This term means a D-ONU that is capable of [802.1] forwarding but cannot do all the encapsulation functions required to be an S-ONU. The B-ONU is a logical definition used by the specification for requirements that apply to all types of B-ONUs. The two types of B-ONUs are the BP-ONU and the BB-ONU.
DPoE Bridge Pluggable ONU (BP-ONU)	This term means a D-ONU that is a B-ONU which is pluggable. Pluggable BP-ONUs include devices such as an SFP-ONU (1G-EPON), SFP+ONU (10G-EPON), or XFP-ONU (10G-EPON).
DPoE Bridge Baseband ONU (BB-ONU)	This term means a D-ONU that is a B-ONU which has a baseband IEEE Ethernet interface. BB-ONUs include those with one or more [802.3] baseband PMDs. (See Section 7.2.6.2 for examples.)
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) or the NID (as in the MEF model).




Figure 3 - D-ONU Types




Figure 4 - DPoE Network Elements

3.2 Other Terms

1G-EPON	EPON as defined in [802.3ah], now part of [802.3].
2G-EPON	EPON as defined in Annex A 2G-EPON System Definition of [DPoE-PHYv2.0]
10G-EPON	EPON as defined in [802.3ah], now part of [802.3].
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 1G-EPON, 2G-EPON, and 10G-EPON collectively
EPON Operations and Maintenance Messaging (OAM)	EPON OAM messaging as defined in [802.3] and [DPoE-OAMv2.0]; Ethernet OAM is not the same as EPON OAM; Ethernet OAM is defined in [802.1ag]
Logical CPE Interface	LCI as defined in [eDOCSIS]
Network Interface Device (NID)	A DEMARC device in DPoE specifications
Service Flow	A unidirectional flow of packets from the upper layer service entity to the RF with pre-defined QoS traffic parameters.
Service Provider	The organization providing Ethernet services.
Subscriber	The organization purchasing and/or using Ethernet services.
TRAN-Trail	A TRAN-trail (see ITU-T Recommendation [G.805]) is a "transport entity" responsible for the transfer of information from the input of a trail termination source to the output of a trail termination sink.

4 ABBREVIATIONS AND ACRONYMS

This specification uses the following abbreviations and acronyms:

ATA	Analog Terminal Adapter
B-ONU	DPoE Bridge ONU
BB-ONU	DPoE Bridge Baseband ONU
BGP	Border Gateway Protocol
BP-ONU	DPoE Bridge Pluggable ONU
BPON	Broadband Passive Optical Network
CBP	Customer Bridge Port
CE	Customer Edge
C-VID	Customer VLAN Identifier
CMCI	Cable Modem CPE Interface
CMIM	Cable Modem Interface Mask
CO	Classifier-ONU
CPE	Customer Premise Equipment
CoS	Class of Service
CS	Classifier-System
CTBH	Cell Tower Backhaul
CWDM	Coarse Wavelength Division Multiplexing
D-ONU	DPoE ONU
DA	Destination Address
DIA	Dedicated Internet Access
DPoE	DOCSIS Provisioning of EPON
DR	Default Router
DSx	Digital Signal (DS1 or DS3)
DVA	Digital Voice Adapter
DWDM	Dense Wavelength Division Multiplexing
eCM	embedded Cable Modem
eDVA	embedded Digital Voice Adapter
ENNI	External Network to Network Interface
EPL	Ethernet Private Line
EP-LAN	Ethernet Private LAN
EPON	Ethernet Passive Optical Network
EP-Tree	Ethernet Virtual Private Tree
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
Gbps	Gigabits per second (as used in the industry)
GBd	Gigabaud
HSD	High Speed Data
IGP	Interior Gateway Protocol
INNI	Internal Network to Network Interface
IP	Internet Protocol
IP(HSD)	High Speed Data (Broadband Internet Access using DOCSIS)
IP-SG	IP Serving Group
IP-VPN	Private IP
iSCSI	Internet Small computer System Interface
I-SID	[802.1ah] I-Component Service Identifier
LCI	Logical CPE Interface
LLDP	Link Layer Discovery Protocol
LLID	Logical Link Identifier
MEF	Metro Ethernet Forum
MEN	Metro Ethernet Network
MI	MEF INNI Interface at a customer premise
MI-SI	MI Service Interface
MLS	Multi-Layer Switching
MN	MEF INNI Interface to operators MEN
MPCP	Multi-Point Control Protocol
MPCPDU	MPCP Data Unit
MPLS	Multiprotocol Label Switching
MSC	Mobile Switching Center
MU	MEF UNI Interface
MU-SI	MU Service Interface
NID	Network Interface Device
NNI	Network to Network Interface
NSI	Network Systems 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
PCP	Priority Control Point
PCS	Physical Coding Sublayer
PDU	Protocol Data Unit
PE	VPLS Capable Provider Edge
PHY	Physical Layer
PMA	Physical Medium Attachment
PMD	Physical Media Dependent (Sublayer)
PON	Passive Optical Network
PSTN	Public Switched Telephone Network
PW	Pseudowire
QoS	Quality of Service
R	IP Router
RD	Route Distinguisher
RIP	Routing Information Protocol
ROADM	Reconfigurable Optical Add-Drop Multiplexer
RS	Reconciliation Sublayer
S-ONU	DPoE Standalone ONU
S-VID	Service VLAN Identifier
SA	Source Address
SCB	Single Copy Broadcast
SD	Service Delimiter
sDVA	Standalone Digital Voice Adapter
SERDES	Serializer/Deserializer
SF	Service Flow
SFP	Small Form-factor Pluggable
SFP+	Small Form-factor Pluggable Plus (+)
SI	Service Interface
SID	Service Identifier
S-ONU	DPoE Standalone ONU
S-VID	Service VLAN Identifier
TDM	Time Division Multiplexing
TDMA	Time Division Multiple Access
TPID	Tag Protocol Identifier
UNI	User Network Interface
V-UNI	Virtual UNI
vCM	Virtual Cable Modem
VE	VPLS Edge

VID	VLAN Identifier
VLAN	Virtual Local Area Network
VOIP	Voice Over IP
VPLS	Virtual Private LAN Service
VPWS	Virtual Private Wire Service
VSI	Virtual Switch Instance
VSI-ID	VSI Identifier
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. The DPoE specifications do not attempt to identify the size of the markets, explain the business interests, nor promote one or more services above any others.

This specification does not explain 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 (e.g., 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 (Metro Ethernet as defined by the 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.

As used in this specification, a service is a set of functions that the operator delivers to a customer which are integrated and managed in such a way that the customer can use the service without the customer having to manage the service. In this specification, an application is a use of a service that is not only determined by the customer, but which is managed or operated by the customer.

Metro Ethernet services, as described by the MEF, are services an operator provides to a customer to transport Ethernet frames associated with that customer's applications. The operation of IP by such a customer, for example, is an application running on that service.

IP(HSD) is a service an operator provides to a customer that provides Internet access for that customer's applications. Operations such as file transfer, web browsing, or streaming video by such a customer are applications running on that service.

In this context, DPoE specifications can be used to provide two basic services:

- Metro Ethernet service (as described by the MEF)
- IP(HSD) (also known as Internet Access)

In addition to these services, both Metro Ethernet services and IP(HSD) services can be used as Ethernet and IP transports, respectively, to transport IP packets or frames from either the DPoE System itself or other network elements and systems within the operator's network, to provide additional services. Because IP can always be transported over a Metro Ethernet service, operators have the option of implementing transport for IP services with IP(HSD) transport or IP over Metro Ethernet services.

DPoE Networks can be used to provide nearly all IP, Ethernet, or IP/Ethernet services either as the service delivery platform or as a transport between the service delivery platform and a demarcation device or D-ONU which provides some portion of the service delivery or at least demarcation of the service delivery.

IP can be used to deliver some storage network services, VoIP, or TDM circuit emulation. T1 or T3 circuit emulation can be used to provide private line PSTN, IP-PBX, cell tower backhaul, and other services.

5.2 Service Requirements

5.2.1 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.

5.2.2 Private Data Services

Private data services are services that are privately used by one or more customers between designated customer premises. These services include:

- Ethernet Service
 - E-Line
 - Ethernet Private Line (EPL)
 - Ethernet Virtual Private Line (EVPL)
 - E-Tree
 - Ethernet Private Tree (EP-Tree)
 - Ethernet Virtual Private Tree (EVP-Tree)
 - E-LAN
 - Ethernet Private LAN (EP-LAN)
 - Ethernet Virtual Private LAN (EVP-LAN)
- Private IP (IP-VPN)

5.2.3 IP (Public IP/Internet)

- IP(HSD) Fixed IP/Internet Service (CMTS IP Routing)
 - Single fixed IPv4 or IPv6 address/MAC
 - Multiple IPv4/IPv4 or IPv6 addresses/MAC
- Internet Address Learning Service
 - RIP-based address learning
 - Border Gateway Protocol (BGP)-based address learning

5.2.4 Voice Services

- May include IP-PBX services (described below) but categorized as voice
- Single line telephony
- Multi-line telephony

5.2.5 Vertical Markets

Any of the above services can be offered for specific vertical markets such as Cell Tower Backhaul (CTBH), which uses any one or more of T1 or T3 circuit emulation, Metro Ethernet, and Internet Access services.

IP multicast services within Metro Ethernet services are not explicitly considered in this version of the specifications. Although IP multicast services can be operated within any Metro Ethernet 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 within Metro Ethernet 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 within Metro Ethernet services.

5.3 Single Tenant Businesses

Single tenant businesses can be served with a dedicated ONU. Each business customer may have one or more applications. A 'service' in DPoE specifications is the network service provided on a single UNI port, such as CMCI for IP(HSD), MU or MI for MEF, or an LCI for an embedded Service/Application Functional Entity (eSAFE).

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. A modular approach can be useful to provide maximum flexibility in meeting the requirements of both single tenant and multiple tenant cases. Such a modular approach could include an S-ONU and a separate DEMARC device or a B-ONU plugged into a DEMARC device.

5.5 DPoE Fully Automated and Partially Automated Services

5.5.1 Private Data Services

5.5.1.1 Ethernet Service

EPON provides 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 to a MEF UNI terminated at an IP router (physical or sub-interface).

Solutions 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 requires a small number of LLIDs. For a multi-tenant solution, it might be necessary to use a D-ONU that has more ports to accommodate a larger number of customers. In that case, the number of required LLIDs for the multi-tenant ONU is expected to be significantly larger.

5.5.1.2 Private IP

IP service can be constructed with EPON by building an Ethernet service from the ONU Ethernet interface to an Ethernet port physically connected to an IP router or to a virtual port logically connected to an IP router 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, Interior Gateway Protocol (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.

5.5.2 Internet Access Service

Internet access service can be offered by two different methods. DPoE specifications describe IP(HSD) service functionally equivalent to that described in DOCSIS, but with higher data rates. The second method is to use Metro Ethernet service, provisioned using DPoE specifications, as a transport and operate IP routing functions for the service on separate operator-managed platforms. This is what operators have historically called Dedicated Internet Access (DIA). Although these can be made to be functionally equivalent for customers, there remains a significant difference between IP over Metro Ethernet transport and IP(HSD). The former requires additional (beyond DOCSIS) provisioning systems and process. The latter can provide a fully automated provisioning process for DPoE Networks that is substantially the same process used for DOCSIS. The resulting savings in engineering and operations is one of the primary objectives of the DPoE specifications effort.

5.5.2.1 Static Address Service

5.5.2.1.1 DPoE IP(HSD) Service

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.

5.5.2.1.2 DPoE Ethernet Transport for DIA Service

The operator can also provide a static IP address service using [DPoE-MEFv2.0] to provide Ethernet service between the customer equipment and another Ethernet interface connected to the operator's IP router.

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.

5.5.2.2 Dynamic Address for Internet Access 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 a DEMARC at the customer premise. It can also be supported by using a B-ONU to provide Metro Ethernet service through a DEMARC.

5.5.3 Internet Transit Service

Internet transit service using BGP can be provided using the same Metro Ethernet services specified in [DPoE-MEFv2.0].

5.5.4 Internet Peering Function

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

5.6 Recommendations for DPoE as a Metro Ethernet or IP Transport

5.6.1 Voice Services

VoIP service other than PacketCable 2.0 can be supported with VoIP or DS1 circuit emulation using DPoE Networks as either an IP transport (with IP(HSD)) or using Metro Ethernet as a transport (with or without IP running over-the-top).

5.6.2 Wholesale Ethernet

The service requirements for wholesale Ethernet are defined by MEF. The requirements for the service (UNI to UNI) are common for all Metro Ethernet services. Wholesale requires the addition of what 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). [MEF 26] is the first of three planned phases that will deliver progressively more complex functionality for ENNI. Carrier interconnection and ENNI are outside of the scope of the access network, and therefore, not defined by DPoE specifications.

5.6.3 T1 or T3 Circuit Emulation

Digital Signal (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 circuit emulation 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, circuit emulation solutions are typically 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 in EPON networks leaves wander as an outstanding problem for circuit emulation solutions. DPoE specifications provide for a synchronization solution in [DPoE-MULPIv2.0].

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.

5.6.4 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 Networks. 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 S-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-MEFv2.0].

5.6.5 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-MEFv2.0].

5.6.6 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-MULPIv2.0] and [DPoE-MEFv2.0].

5.6.7 Storage Network Services

DPoE specifications do not contain any specific requirements or recommendations for storage network services. As with all Ethernet-based services, EPON can only support IEEE standard frame sizes. As an IEEE 802 Ethernet protocol, 1G-EPON, 2G-EPON, and 10G-EPON do not typically transport frames larger than 2000 bytes² as specified by [802.3]. 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. It is expected, though, that new generations of 1G-EPON and 10G-EPON chipsets will support frame sizes larger than 2000 bytes to support mobile backhaul and data center applications. The frame size requirements for DPoE Networks are specified in [DPoE-MULPIv2.0].

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-MEFv2.0].

5.6.8 Cloud Services

Both Metro Ethernet and IP(HSD) services can be used to provide network access to cloud services.

5.6.9 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 an S-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 C₀ reference point for B-ONU implementations. Such a management connection is no different than other EVCs.

For an S-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 D-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 interface where there is a single EVC on a single S interface. In the third model, the operator would construct an over-the-top EVC within an EPL or EVPL through an S 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 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.

² IEEE 802.3 support for 1500 byte payload remains the same in both 1G-EPON, 2G-EPON, and 10G-EPON. The additional 100 bytes (1600 byte total) in 1G-EPON and 2G-EPON or the additional 500 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.

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.

6 ARCHITECTURAL FOUNDATION

6.1 Forwarding

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

6.1.1 Ethernet Virtual Connections

Services constructed with DPoE Networks use 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 Networks.

In [802.1ad], the term S-VID refers to a Service (S) VLAN Identifier, and the term C-VID refers to a Customer (C) VLAN Identifier. Services implemented according to DPoE specifications use the S-VID as an “outer” VLAN identifier and the C-VID as an “inner” VLAN identifier without respect to their typical meaning in [802.1ad].

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

- 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

For some services, the EVC may terminate at a Virtual UNI (V-UNI)³, sub-interface, attachment interface, or similar interface within the hub. For other services, the EVC may traverse a Layer 2 Metro Ethernet Network (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 over a DPoE Network will be an EVC from a DEMARC device to a D-ONU to a DPoE System. From the DPoE System, the EVC trail can be "spliced" into Ethernet transports or "terminated" at an IP router interface for IP transport and services. Because IP routing functionality is an integral part of a DPoE System, the IP service can be provisioned using DPoE specifications.

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 DPoE specifications call a DEMARC), and many types of INNI interfaces at the MN interface. Each service that MEF describes needs to be supported over each of these different models. MEF itself supports [802.1Q], [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. The specifications for Metro Ethernet services are described in [DPoE-MEFv2.0].

6.1.3 Internet Services

The IP(HSD) service is implemented using DPoE specifications in one of two ways. The first implementation method is equivalent to a "bridged" service from the customer premise through a D-ONU interface configured as a CMCI, to a default router (DR) operating on a DPoE System. The second implementation is [eRouter], which operates an IP router within the S-ONU. In this implementation, the bridged service is from the router in the DPoE

³ As described by MEF, a V-UNI operates at an ENNI interface. Whether that ENNI interface operates within a single operator's network that is logically divided into separate operational networks, or literally between two different operators, does not affect the requirements for such a V-UNI because they are the same functional requirements.

System to an LCI interface on an eRouter (which is a type of eSAFE) in an S-ONU that is operating as an eDOCSIS device.

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 services to the MAC and MAC sub-layers in order to provide consistent Quality of Service based on the TDM capabilities of EPON. In DPoE version 2.0 specifications, an LLID may carry more than one Metro Ethernet service as described in [DPoE-MEFv2.0] and [DPoE-MULPIv2.0].

6.3 Multipoint Architecture

One fundamental difference between EPON, DOCSIS, and Ethernet 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 optical distribution network. EPON supports forward broadcast, multicast, and unicast when using the MAC. Upstream transmissions in EPON are not broadcast to all ONUs. In the upstream direction, optical signals from individual ONUs are passively combined into a single return path and optical combiners are highly unidirectional. The EPON TDMA implementation features a master scheduler in the OLT to time division multiplex upstream burst 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. 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.

EPON's MPCP, in combination with the service segregation 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 services.

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 D-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 D-ONU from the OSS passes through the DPoE System, it is possible for the DPoE System to perform the IP layer functions of a CM. DPoE Systems instantiate a virtual Cable Modem (vCM) for each registered D-ONU. The vCM handles all the OAMP functions for DOCSIS as described in [DPoE-MULPIv2.0] and [DPoE-OSSIV2.0]. The vCM can proxy requests, signaling, and messages to the D-ONU using EPON OAM messages defined in [DPoE-OAMv2.0].

The vCM model applies only to the D-ONU and not to other embedded components that may be present in the D-ONU. eDOCSIS devices (if present) use an IP stack on the D-ONU for all other DOCSIS services beyond CMCI-based services. This includes Embedded Digital Voice Adapter (eDVA), eRouter, or any other eSAFE subsystem within a D-ONU.

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 IP address is used only within the operator's network.

The concept of a DOCSIS Service Flow (SF) is maintained in the DPoE architecture. Services transmitted over a DPoE Network are classified and assigned to an SF. An LLID is the mechanism for guaranteed delivery of SFs over

the DPoE Network. One or more SFs may be aggregated and transmitted using a single LLID. An explanation of SF aggregation is provided in [DPoE-MULPIv2.0].

Since each service has separate requirements, 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. Like the SID scheduler in DOCSIS, 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.

6.5 eDOCSIS

As shown in Figure 5, an eDOCSIS device defined in [eDOCSIS] consists of an embedded DOCSIS cable modem (eCM) and one or more eSAFEs. An eDOCSIS device may also have one or more physically exposed interfaces. DPoE specifications allow for the optional support of [eDOCSIS]⁴. An S-ONU that includes an eRouter, eDVA, or other eSAFEs is considered an eDOCSIS device.




Figure 5 - eDOCSIS Reference Model (from Figure 5-1 of [eDOCSIS])

The DPoE architecture calls for the IP management functionality of a CM to reside in the DPoE System. Thus, similar to a CM without an eSAFE, and notwithstanding the downstream and upstream PHY, the functionality of the eCM in the eDOCSIS device is correspondingly split between the DPoE System and the S-ONU. The eSAFE functionality remains unchanged and is described in [eDOCSIS]. This arrangement is shown conceptually in Figure 6. Note the ePTA has no downstream external interface.

⁴ The SLED functionality defined in [eDOCSIS] is not explicitly supported in DPoE specifications.




Figure 6 - eCM Functionality in an eDOCSIS Device Split Between DPoE System and S-ONU

Consistent with the DPoE architecture, the portion of the eCM that resides in the DPoE System is identical to the vCM, plus additional requirements levied on the eCM in [eDOCSIS]. The remaining functionality of the eDOCSIS device, which includes classification, forwarding, QoS enforcement, and eSAFE functionality, is implemented in the S-ONU, as shown in Figure 7. The vCM that is created on behalf of an eDOCSIS S-ONU must satisfy the IP management requirements defined in [eDOCSIS], while the eDOCSIS S-ONU must satisfy the remaining requirements in [eDOCSIS].




Figure 7 - S-ONU that is an eDOCSIS Device, with vCM in DPoE System

6.6 EPON OAM

As described above, the control path for communication between the DPoE System and D-ONUs are EPON OAM messages and OAM Extensions as defined in [802.3] and [DPoE-OAMv2.0]. EPON OAM messages are datagrams available only on the PON and sent between the OLT and ONU. EPON OAM messages are not forwarded beyond the OLT or beyond the ONU. The capabilities of EPON OAM and the related extensions defined in the DPoE specifications include ONU configuration, alarm and event management, and statistics gathering. EPON OAM is distinct from Ethernet service OAM, such as [802.1ag] or MEF OAM as defined in [MEF 21].

7 REQUIREMENTS

7.1 Introduction

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

One very fundamental difference between DOCSIS and DPoE Networks is that the least common denominator of services in DOCSIS is IP, whereas the least common denominator of services 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 Elements

A DPoE Network element is a device for which DPoE specifications provide requirements. The DPoE Network elements and their corresponding abbreviations are:

Table 3 - DPoE Network Element Abbreviations

DPoE Network Element	Abbreviation
DPoE System	
DPoE ONU	D-ONU
DPoE Standalone ONU	S-ONU
DPoE Bridge ONU	B-ONU
DPoE Bridge Pluggable ONU	BP-ONU
DPoE Bridge Baseband ONU	BB-ONU

7.2.2 DPoE Network Definition

The term DPoE Network means the entire network described in Figure 2 from the D or MN interface to the LCI, CMCI, S, MI, or MU interface. This includes one DPoE System, one PON, and all connected D-ONUs.

7.2.3 DPoE System Definition

The DPoE System is analogous to the CMTS in DOCSIS networks. The DPoE System is the control point for the DPoE Network and is responsible for allocating upstream bandwidth to the connected D-ONUs.

The DPoE System refers to one or more physical devices that are required to mimic the CMTS. It provides 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-IPNEv2.0]. The DPoE System is depicted in Figure 4. While the DPoE System may consist of more than a single networked device, it MUST appear to the OSS as a single device.

7.2.4 DPoE ONU Definition

This term means a DPoE-capable ONU (D-ONU) that complies with all the DPoE specifications. There are two logical types of D-ONUs. These are the DPoE Standalone ONU (S-ONU) and the DPoE Bridge ONU (B-ONU). The D-ONU is a logical definition used by the specification for requirements that apply to all types of ONUs.

7.2.5 DPoE Standalone ONU Definition

The DPoE Standalone ONU (S-ONU) is intended to support both IP(HSD) services and Metro Ethernet services without the need for an external device, such as a DEMARC device. Although an S-ONU does not require a DEMARC to provide a UNI interface such as MU, an S-ONU can be optionally used with a DEMARC connected to the UNI.

An S-ONU MAY contain more than one Ethernet port and support a variety of [802.3] physical media. If an S-ONU has more than one Ethernet port, it MAY also support a variety of bridging and switching functions. An S-ONU MAY support internal LCI interfaces. Specifically, an S-ONU MAY include a PacketCable 2.0 eDVA, a DOCSIS 3.0 eRouter, DSG, 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 an eDOCSIS device interface providing an architecture equivalent to the eDOCSIS architecture. An S-ONU MUST use a dedicated LLID for every service flow associated with an LCI.

7.2.6 DPoE Bridge ONU Definition

The term DPoE Bridge ONU (B-ONU) refers to a D-ONU that MUST have at least one S interface that can be operated in "Transport Mode" as described in [DPoE-MEFv2.0]. One example of a B-ONU is an SFP-ONU. Another example of a B-ONU is a simple ONU based on an EPON chipset with one or possibly two Ethernet ports, but with no switching, forwarding, nor encapsulation capabilities beyond those of the EPON chipset.

7.2.6.1 DPoE Bridge Pluggable ONU Definition

The term DPoE Bridge Pluggable ONU (BP-ONU) refers to a sub-type of B-ONU that is a modular EPON ONU and EPON transceiver in a single package or device that can be plugged into a DEMARC device using an interface such as those described below in Table 4. A BP-ONU typically receives power from the DEMARC. Examples of BP-ONUs include:

Table 4 - Examples of BP-ONUs

BP-ONU Description	BP-ONU Abbreviation
Small Form-factor Pluggable transceiver plus 1G-EPON ONU	SFP-ONU
Small Form-factor Pluggable transceiver plus 10G-EPON ONU	SFP+ONU
X Form-factor Pluggable transceiver plus 10G-EPON ONU	XFP-ONU

The BP-ONUs provided here are examples only. Other form factors of BP-ONU are possible and permissible. BP-ONUs are expected to meet all requirements written for D-ONUs and B-ONUs.

7.2.6.2 DPoE Bridge Baseband ONU Definition

The term DPoE Bridge Baseband ONU (BB-ONU) refers to a type of B-ONU that is independently powered from other elements and which uses an [802.3] baseband PMD. While the complete list is provided in [802.3], examples of such PMDs include 100BASE-T, 1000BASE-T, 1000BASE-FX, 10GBASE-LR/SR, etc.

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 Network elements – that is, the inputs and outputs between the OSS and the DPoE System, between the DPoE System and a D-ONU, and between the D-ONU and various types of interfaces to CPE, eDOCSIS devices, and DEMARC devices. Since S-ONUs can contain

embedded subsystems (eDOCSIS devices) that are often built by OEMs other than the ONU vendors, the DPoE specifications include a set of interfaces within the D-ONU called the S interface.

There are two types of DPoE interfaces. These are external interfaces and internal interfaces.

7.3.1.1 External Interfaces

External interfaces (referred to only as "interfaces") are physical interfaces that can be connected to other DPoE Network elements or external devices, and can be tested. Interfaces are accessible for direct testing.

7.3.1.2 Internal Interfaces

Internal interfaces are those that can be described and specified, but which cannot be directly tested. An example of an internal interface is the MN_I interface. The function of MN_I can only be indirectly tested by configuring the DPoE System and using it in a DPoE Network that is also connected to an IP network with VPLS. Although such an internal interface (like a reference point) cannot be directly tested, it is further distinguished from a reference point in that the actions of a DPoE Network element at an internal interface can be indirectly observed. In the example of MN_I, the output at the D interface can be tested to show conformance with specifications for MN_I.

7.3.1.3 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
- OAMP functionality for IP and Ethernet services
- VPLS or VPWS PWE encapsulated traffic for Metro Ethernet services

Each D interface on the DPoE System must support IP routing and forwarding as described in [DPoE-IPNEv2.0]. The D interface is also the interface for interoperability between the operator's OSS and DPoE Systems.

7.3.1.4 TU Interface

The TU interface is the interface between a DPoE System and one or more D-ONUs. Any DPoE System is intended to interoperate with any D-ONU. Each TU interface supports at least one TUL.

7.3.1.5 TUL Interface

TUL is a virtual interface used to describe a MAC domain on a DPoE Network. TUL and MAC domains are described in [DPoE-MULPIv2.0] and are used in [DPoE-OSSIV2.0], and [DPoE-IPNEv2.0].

7.3.1.6 S Interface

The S interface represents the customer interfaces in the DPoE Network on D-ONUs. An S interface is present on D-ONUs with external PHY interfaces operating as CMCI, MI or MU. An S interface MAY also be hardwired to an eDOCSIS device (e.g., an LCI) or as an internal control interface (such as an [802.3] GMII or XGMII interface implemented in the form of a SERDES). Specifications for provisioning ports to operate as MI or MU are in [DPoE-MEFv2.0]. Specifications for provisioning ports to operate as CMCI are in [DPoE-MULPIv2.0].

7.3.1.7 LCI Interface

The LCI interface is equivalent to an eDOCSIS Logical CPE Interface (LCI) as described in [eDOCSIS]. An S interface connected to an eDOCSIS device such as an eDVA, eRouter, or DSG is always an LCI interface.

7.3.2 DPoE Reference Interfaces and Reference Points

Reference interfaces are interfaces used on DPoE Network elements for which the requirements are identified by reference to another standard, recommendation, or specification other than DPoE or DOCSIS.

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 Interface

The virtual 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. The requirements for MN are specified in [DPoE-IPNEv2.0]. 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 MN_E Reference Interface

The MN_E (MEF INNI External) interface is substitute for the MN reference interface from DPoE version 1.0 specifications. The MN_E 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.

7.3.2.3 MN_I Reference Point

The MN_I reference point is used to describe the virtual interface between an OLT and a VPLS VSI. In particular, it is used to describe the requirements for stitching VPLS VSIs to DPoE System and OLT [802.1] components such as [802.1Q] bridge groups, [802.1ad] S-VLAN or C-VLAN (S-component or C-component), or [802.1ah] I-BEB (I-component) or B-BEB (B-component) backbone edge bridges. The DPoE System stitches VPLS and VPWS transport and forwarding for Metro Ethernet services between the D interface and the MN_I reference interface. Converged IP(HSD), management [DPoE-IPNEv2.0], [DPoE-OSSv2.0], and [DPoE-SECv2.0]), and IP/MPLS/VPLS or IP/MPLS/VPWS traffic is carried on the D interface.

7.3.2.4 S Reference Point

The S reference points are used to describe either virtual (logical) or real (external) Ethernet ports on D-ONUs. The S reference point is used to provide a logical means for common requirements for all D-ONU interfaces. Requirements for the S reference points are the same as those for the S interface. An S reference point performs the same function as an S interface, but this specification cannot formally specify requirements for that reference point because the element is not fully specified by DPoE.

7.3.2.5 CS Reference Point

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

7.3.2.6 CO Reference Point

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

7.3.2.7 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 [802.3]/[802.1]. 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).

7.3.2.8 MU Reference Point

An interface that provides the MEF UNI on a D-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 specifications. The MU reference point is used to describe the MEF User to Network Interface (UNI) on S-ONUs or DEMARC devices.

7.4 Architectural Requirements

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

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

All DPoE Network elements require [802.3ah], and, if applicable, [802.1], and IETF protocol support (as defined in [DPoE-IPNEv2.0] and [DPoE-OSSv2.0]). The architecture does not require any interoperability between subsystems 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, IP transport interfaces D, and the MEF Ethernet interfaces MN, MI, and MU to be interoperable with other systems.

7.4.1 Ethernet Service Based Architecture

The Metro Ethernet service model for a DPoE Network uses tag-based 802.1 Ethernet switching and forwarding, augmented with DOCSIS-based provisioning and configuration automation mechanisms. The tag-based switching and forwarding uses a combination of [802.1ad] or [802.1ah] tags to establish a data path between a D-ONU and DPoE System. Frames belonging to the same service share the same [802.1ad] or [802.1ah] tags and follow the same path through the DPoE Network, as configured by the operator. Each Metro Ethernet service is granted a dedicated [802.1ad] or [802.1ah] EVC spanning between the D-ONU and DPoE System. IP(HSD) is a DPoE service that transports IP packets between D-ONU and DPoE System to an IP router that is part of the DPoE System. Forwarding decisions for IP(HSD) frames are based on a combination of dynamic MAC learning or static MAC address provisioning, supporting the [802.1Q] bridged service model.

Various modes of operation are supported in the DPoE Network, where the D-ONUs and DPoE System may (i) perform encapsulation or de-encapsulation (so-called Encapsulation Mode) of subscriber frames by adding or stripping configured [802.1ad] or [802.1ah] tags or (ii) relay subscriber frames transparently (so-called Transport Mode) without making changes to already existing [802.1ad] or [802.1ah] tags. The D-ONU and DPoE System may be configured to perform either one of these modes independently to achieve the target service connectivity across the DPoE Network. The decision on which mode to configure individual DPoE Network elements is determined by the operator based on the vCM configuration.

Details about the MEF services, their configuration and provisioning mechanisms, as well as types of supported services can be found in [DPoE-MEFv2.0]. Details about the IP(HSD) services, their configuration and provisioning mechanisms, can be found in this specification and [DPoE-MULPIv2.0].

7.5 Service Model

DPoE services are based on the operator requirement to deliver a high QoS for each service. The DPoE specifications utilize the LLID by associating either multiplexed or individual services with those LLIDs.

7.5.1 Service Multiplexing Model

In order to provide strict scheduler control for each service, a D-ONU MUST support one LLID for each provisioned service. A D-ONU MAY aggregate traffic from two or more Metro Ethernet services received from MU or MI interfaces onto the same LLID.

Specifications for provisioning ports for MI or MU services can be found in [DPoE-MEFv2.0]. Specifications for provisioning ports for CMCI or LCI service can be found in [DPoE-MULPIv2.0]. DPoE Network forwarding and configuration for DOCSIS modes are described in [DPoE-MULPIv2.0]. DPoE Network forwarding and configuration for Metro Ethernet services are described in [DPoE-MEFv2.0].

Management traffic for embedded subsystems or external devices may be carried in a separate EVC. For example, dedicated management for an eSAFE would be directed to that eSAFE device via proper configuration of the classifiers on the DPoE System and the DPoE ONU. Operators may choose to provision a dedicated LLID for such eSAFE management traffic, if bandwidth requirements imposed by such traffic justify allocation of a dedicated LLID.

An LCI interface is an internal interface that connects only to an eSAFE such as, e.g., eRouter, eDVA, or other eSAFEs. Each eSAFE uses a single LCI interface and the LCI interface is also a type of S interface. Customer-facing eSAFE ports are not included in the port count for the minimum number of required LLIDs. For example, an eSAFE with two customer-facing ports is counted as a single S interface for minimum LLID requirements.

A D-ONU MUST support at least eight (8) LLIDs. A D-ONU MUST support at least one (1) LLID for each S interface supported by the D-ONU. A D-ONU SHOULD support at least two (2) LLIDs for each S interface supported by the D-ONU. A D-ONU SHOULD support at least one (1) LLID for each customer-facing eSAFE ports.

Consider two examples of D-ONU configurations with multiple types of interfaces.

A D-ONU with four (4) S interfaces, with one of those S interfaces connected to an eRouter that provides four (4) customer-facing eSAFE ports, is required to have at least eight (8) LLIDs. Ideally, the D-ONU should provide 12 LLIDs (two (2) LLIDs x four (4) S interfaces, plus one (1) LLID x four (4) customer-facing eSAFE ports).

A D-ONU with four (4) S interfaces, with one of those S interfaces connected to an eDVA that provides one (1) customer-facing eSAFE port is required to have at least eight (8) LLIDs. Ideally, the D-ONU should provide nine (9) LLIDs (two (2) LLIDs x four (4) S interfaces, plus one (1) LLID x one (1) customer-facing eSAFE port).

7.5.2 Classification Reference Points

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

7.5.2.1 Upstream Classification Reference Point

D-ONUs MUST perform all upstream traffic encapsulation, tagging, marking, and transcoding of frames at or before the C_O reference point. Upstream traffic encapsulation and forwarding are described in [DPoE-MULPIv2.0] for IP(HSD) and [DPoE-MEFv2.0] for Metro Ethernet services. Classification and scheduling for all services are described in [DPoE-MULPIv2.0]. By definition, all upstream classification is performed at or before the C_O reference point. Some traffic encapsulation, tagging, marking, and transcoding of frames or packets is performed on the DPoE System, but always prior to the D or MN interface.

7.5.2.2 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-MULPIv2.0] for IP(HSD) and [DPoE-MEFv2.0] for Metro Ethernet services. Classification and scheduling for all services are described in [DPoE-MULPIv2.0]. By definition, all downstream classification is performed at the C_s reference point.

7.6 IP(HSD) Forwarding Model

The IP(HSD) forwarding model is described in this specification and also in [DPoE-MULPIv2.0]. The DPoE System MUST include a fully functional IP router to meet all the requirements in [DPoE-IPNEv2.0]. These are the same requirements that operators have for DOCSIS-based CMTS systems.

IP(HSD) frame forwarding is based on [802.1Q] bridging. The DPoE ONU and the DPoE System (in combination) implement the MAC learning function, where the source MAC address of the received frame is recorded in a local MAC address table, creating an association with the given UNI port / LLID and the specific CPE MAC address.

In the downstream direction, data forwarding through the DPoE System is a combination of transparent bridging and network layer forwarding. After network layer processing completes, the DPoE System uses the aforementioned MAC address table to map the destination MAC address of a frame to the appropriate LLID for transmission over the ODN. Once the frame is received by the D-ONU, the D-ONU may consult its MAC address table to forward the frame out the appropriate interface.

Specific requirements for DPoE Systems and D-ONU are provided in [DPoE-MULPIv2.0].

7.7 L2HSD Forwarding Model

L2HSD is a simple VLAN-based L2 bridge service instantiated per Serving Group using the same constructs as IP(HSD) (i.e., TLVs and DPoE System configuration parameters as described in [DPoE-IPNEv2.0]). The DPoE System identifies an L2HSD service by an S-VLAN definition in a serving group configuration.

L2HSD is similar to L2VPN functionally and both may be simultaneously instantiated within a single DPoE system. However, S-VLANs are not shared between L2HSD and L2VPN services. Basic multipoint operation is expanded to include user-to-user communication restrictions similar to those in VLAN-rooted multipoint configurations.

7.8 eDOCSIS

DOCSIS 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 LCI.

7.8.1 eDOCSIS on S-ONU

Any S-ONU that implements a DOCSIS eSAFE such as eRouter, eSTB, or eDVA is an eDOCSIS device and MUST support [eDOCSIS]. An S-ONU that is an eDOCSIS device MUST have an IP stack on the S-ONU for each eSAFE. This includes eDVA, eRouter, or any other eSAFE subsystem within an S-ONU. External CPE devices such as a DVA uses IP stacks on each respective device. The vCM model is only used for DOCSIS CM management functions.

An S-ONU that is an eDOCSIS device MUST implement the "SF Classifier" and "MAC Bridge and Filter" functions of the eCM.

7.8.2 eSAFE Discovery

A vCM MUST use [DPoE-OAMv2.0] PDUs to query an S-ONU for embedded components. When queried by the DPoE System, an S-ONU that is an eDOCSIS device MUST use [DPoE-OAMv2.0] PDUs to indicate embedded components to the vCM. The indicated embedded components are then used to populate subsequent DHCP messages sent by the vCM to the OSS, as described below.

7.8.3 vCM and eCM

DPoE Systems MUST 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 vCM MUST use DHCP Option 43, detailed in [eDOCSIS], to report eSAFEs present in an S-ONU. When a vCM receives and proxies a software image for an S-ONU that is an eDOCSIS device, the S-ONU software image MUST be a monolithic software image for the eDOCSIS device, as described in [eDOCSIS].

DPoE Systems MUST use [DPoE-OAMv2.0] PDUs for all required eCM functions described in [eDOCSIS]. S-ONUs MUST use [DPoE-OAMv2.0] PDUs for all required eCM functions described in [eDOCSIS].

7.8.4 eDOCSIS Applicability

An eDOCSIS device consists of an embedded DOCSIS cable modem (eCM) and one or more eSAFEs. The eCM is connected to an eSAFE via a Logical CPE Interface (LCI). An eDOCSIS device may also have one or more physically exposed interfaces (CMCI ports). The eDOCSIS model defines a set of eSAFEs (e.g., eMTA, eSTB, eRouter,) and also allows for "other eSAFEs" with "other interfaces" (as described in [eDOCSIS]). The DPoE specifications continue to use the eDOCSIS model and allow for the support of "other eSAFEs" within an eDOCSIS device. An example might be an embedded MoCA or WiFi device which acts as a PHY bridge. eDOCSIS also specifically includes interfaces from the internal MAC bridge directly to Physical CPE interfaces.

With each LCI connection in an S-ONU (as shown in Figure 8), there is a uniform model for provisioning eSAFEs. For external devices connected to a CMCI port (and not an LCI), the OSS is not aware of the QoS that may be necessary for that external device. Nevertheless, if a particular S-ONU provides dedicated ports for specific external devices, operators may provision specific QoS for the interface connected to the external device.




Figure 8 - DPoE Specifications Support Both Embedded SAFEs and External CPE

7.9 Power Saving

Power conservation and reduction of the carbon footprint of access networks is globally recognized as one of the technical targets for the optical access networks. The objectives of the power-saving mechanisms are to reduce ecological impact, reduce operating cost, and extend battery backup time (if supported by the given product), while minimizing any degradation of network performance to maintain the configured SLA. The scope of the power saving mechanism defined for DPoE covers the operation of the ONU, DPoE System, as well as the interaction between the DPoE System and the ONU (control protocol).

There are no requirements for the performance of the power saving mechanism (i.e., number of watts saved) when enabled on the ONU and the DPoE System.

In the scope of DPoE specifications, the power saving mechanism is defined via reference to [1904.1A], with details included in [DPoE-MULPIv2.0] and [DPoE-OAMv2.0] specifications. The objectives of the power saving mechanism are specified in [1904.1A], subclause 10.4.

7.10 Optical Link Protection

Optical link protection provides redundancy for the data path between the DPoE System and ONUs connected to one of the TU interfaces on the DPoE System. Two types of optical link protection schemes are supported in [1904.1A], i.e., namely, a trunk protection (see [1904.1A], subclause 9.3.3) and a tree protection (see [1904.1A], subclause 9.3.4), each addressing a different application space and providing different types of functionality. The scope of the optical link protection mechanisms defined in DPoE specifications covers the operation of the ONU, DPoE System, as well as the interaction between the DPoE System and the ONU (control protocol).

There are no requirements for the performance of the optical link protection mechanism (i.e., maximum number of lost packets, maximum down time, etc.) when enabled on the ONU and the DPoE System.

In the scope of DPoE specifications, the optical link protection is defined via reference to [1904.1A], with details included in [DPoE-MULPIv2.0] and [DPoE-OAMv2.0] specifications. Device architecture is described in [1904.1A], subclause 9.3.2, with trunk and tree protection mechanisms defined in [1904.1A], subclause 9.3.3 and 9.3.4, respectively.

8 EVPL, E-LAN, AND E-TREE SERVICES

In DPoE 1.0 specifications, a single MU equated to a single EPL service instance end-point. In this version of the specifications, with an EVPL service, a given MU interface can have multiple service instances attached to it, commonly conceptualized as sub-interfaces or C-tagged service interfaces [802.1Q].

For the purpose of this document, these specific service interfaces on an MU will be called MU Service Interfaces (MU-SI). Specific service interfaces on an MI will be referenced as MI Service Interfaces (MI-SIs.) An MI-SI or MU-SI can include a pool of matched VIDs (as in bundling or all-to-one bundling) or just a single matched VID. An MI-SI or MU-SI may also include a set of matched CoS identifiers (such as Ethernet Priority Control Point (PCP) values, IP Precedence values, DSCP values, or Traffic Class values.) The generic term SI will be used to indicate a service instance match based strictly on the classification rules in a transport-mode service or based on the classification rules and [L2VPN] encapsulation, regardless of whether the service interface exists on an MI or an MU.




Figure 9 - ELAN-MPLS DPoE Model

To provision these advanced services, an operator will provision the classifiers, service flows, and encapsulation NSI as they do with all MEF services. These values will configure the TRAN-trail (see Figure 9) between the D-ONU and the virtual switch instance (VSI) within the Router/PE. A new CM configuration object (the service delimiter) will be utilized to associate the TRAN-trail via its [802.1Q] encapsulations on the MN_I interface, the interface where all the TRAN-trails must traverse to reach the router/PE inside the DPoE System.

Additional configuration elements (defined later within this document) will be utilized to provision the specific local service elements. These can be either pseudowire configuration or VPLS VSI configuration. The pseudowire and VPLS VSI configuration can either be completely provisioned via the CM configuration file or can be provisioned via a reference to a pseudowire or VSI configured manually on the DPoE System. The subsequent text describes these varying operations in greater detail.

8.1 E-Line Service

This specification supports the automation and provisioning of an E-Line where both SIs exist on D-ONUs on different DPoE Systems, where both SIs exist on D-ONUs connected to the same DPoE Network, as well as where both SIs exist on the same D-ONU. In all aforementioned cases, the service frames must traverse the connected DPoE System before being forwarded from one SI to another SI. In other words, this specification does not support the direct forwarding between two SIs participating in the same E-Line on the same D-ONU without being forwarded through the connected DPoE System.

In addition, this specification supports the automation and provisioning of one-half of an E-Line in the case where one end of the E-Line exists on a D-ONU MU-SI or MI-SI and the other end of the E-Line exists on a MEF UNI on some other non-DPoE device. In this case, this specification would support the automated provisioning of the MU-SI and the supporting service-layers that connects to the D-ONU, and the operator would be responsible for configuring the UNI and supporting service-layers that connect to the non-DPoE-provisioned device via a separate mechanism (i.e., manual configuration, separate dynamic configuration process, etc.).

8.2 E-LAN and E-Tree Services

This specification supports the automation and provisioning of an E-LAN and E-Tree where the SIs being automated exist on a D-ONU. This specification also supports the automation and provisioning of multiple SIs provisioned on the same D-ONU, different D-ONUs on the same DPoE System, as well as D-ONUs on different DPoE Systems. As a result, multiple TRAN-trails can exist on the same DPoE System and participate in the same E-LAN or service instance.

In addition, this specification supports the automation and provisioning of the SIs participating in an E-LAN or E-Tree that exist on a D-ONU in conjunction with some alternately-provisioned SIs whose UNIs exist on some non-DPoE-provisioned subject to the restrictions at the end of Section 8.1.

8.3 Supported Forwarding Models

This specification supports two general learning and forwarding models: (1) Distributed – the [802.1] bridging/learning occurs locally on the DPoE System where the service endpoint is provisioned and (2) Centralized – the [802.1] bridging/learning occurs on some remote system (i.e., not on the DPoE System where the service interface is provisioned). In the case of the centralized bridging model, the DPoE System still participates in forwarding behavior between the local attachment circuit and the pseudowire (PW) forwarder. To accommodate this, the bridging/learning behavior is provisioned along with the service instance.

As a result, the bridge network topology could be provisioned to be fully centralized, fully distributed, or it could be a hybrid of the two, based on operator preference.

8.3.1 Fully Centralized

In the fully centralized model, the MAC learning and forwarding occurs within a "centralized" bridge or Virtual Switching Instance (VSI). Thus, per operator provisioning, each DPoE System would establish a spoke pseudowire or VPWS per participating SI to the centralized VSI. The centralized VSI would forward the frames based on standard [802.1Q] bridging (unqualified learning via 802.1d; or qualified learning via 802.1d + [802.1Q]) behavior to the appropriate pseudowire spoke, VPWS, or local attachment circuit.

8.3.2 Fully Distributed

In the fully distributed model, each DPoE System with a participating SI in an E-LAN or E-TREE service instance would contain and operate its own local VSI. The DPoE System would forward frames within that VSI based on standard [802.1Q] bridging (unqualified learning via 802.1d; or qualified learning via 802.1d + [802.1Q]) behavior to the appropriate pseudowire spoke, VPWS, or local attachment circuit.

8.3.3 Hybrid Centralized-Distributed

Because the type of switching is based on local provisioning, per SI, an operator could provision a hybrid service among multiple DPoE Systems, or among some set of DPoE Systems and other PEs. An operator could provision some SIs as VPWS spokes to a centralized VSI in the cloud and other SIs to a distributed (local) VSI on the DPoE System. Via [RFC 6074]-based implementations, it would then be possible for the local VSI to construct a topological relationship to the VSI in the cloud such that an end-to-end service instance exists for the EVC.

8.4 EVPL, E-LAN, E-Tree Service Provisioning

The general EVPL, E-LAN, and E-Tree service provisioning model is identical to the provisioning model for all other services in that there is some configuration file that the vCM loads when it boots. The vCM then parses the configuration. For part of the configuration, the vCM generates extended OAM messages which configure the behavior of the D-ONU; this includes configuring objects like the downstream classifiers, MESP parameters, etc. For the remainder of the configuration file, the vCM generates local configuration for the DPoE System service instance.

8.4.1 DOCSIS-L2VPN Service Provisioning Model

In the [L2VPN] provisioning model, using applicable TLVs from [L2VPN], an operator can configure an EVPL via VPWS as an ELAN or E-Tree spoke to a centralized forwarding entity using the VPWS to connect to a centralized VSI. The pertinent TLVs from [L2VPN] are: VPN-ID, MPLS Peer, Pseudowire ID, and Pseudowire Type. There are a number of TLVs required and defined in [DPoE-MEFv2.0] to provide enhanced capabilities above that already defined in [L2VPN]. The set of TLVs includes the Backup MPLS Peer, Backup Pseudowire ID, PW-Class, and the Service Delimiter (SD) TLV.

An important distinction between the MPLS service provisioning in [L2VPN] for DOCSIS from the similar capabilities described in this text is that in the DPoE specifications, both the TRAN-trail (classifier, service-flow, encapsulation [802.1Q] tags) and the pseudowire configuration elements need to be provisioned. In [L2VPN] either [802.1Q] encapsulation is provisioned or MPLS-based pseudowires are provisioned.

A brief description of each of the new TLVs is provided here, though the detailed description and interaction between the TLVs is defined in [DPoE-MEFv2.0]. The extensions to the Encapsulation NSI TLV (43.5.2) are shown below in Figure 10.




Figure 10 - Encapsulation NSI Extension

- Service Delimiter: the selector-bytes used by the Bridge Forwarder to associate the frames from a defined TRAN-trail to the service instance or pseudowire forwarder. The Service Delimiter TLV has a number of sub-options including the SD-C-VID, SD-S-VID, SD-I-SID, and SD-B-VID. This provides the flexibility of the DPoE System VPLS Edge (VE) to associate the proper TRAN-trail to the VSI.
- Backup MPLS Peer: the backup MPLS peer is used to define a second Pseudowire destination IP for use in pseudowire redundancy. The protocol requirements for pseudowire redundancy are included in [DPoE-IPNEv2.0].
- Backup Pseudowire ID: this is the backup MPLS Pseudowire ID also needed to define the second pseudowire endpoint for use in pseudowire redundancy.

- PW-Class: the PW-Class is used to define standard pseudowire behavior via a class locally configured on the DPoE System. The PW-Class is a configuration file reference to an object configured on the DPoE System. The PW-Class includes objects such as pseudowire MTU, service delimiter tag pop or retain actions, and pseudowire redundancy behavior such as primary PW preemption. The complete description of the PW-Class object is included in [DPoE-IPNEv2.0].
- If a PW is configured where both endpoints exist on the same DPoE System, where the DPoE System notes that both endpoints a) have the same VPN-ID, b) have the same MPLS Peer, and c) have the same Pseudowire ID, the DPoE System MUST configure the service as an ethernet cross-connect.

8.4.1.1 Example DOCSIS L2VPN Pseudowire Service Configuration

To provision an MPLS Pseudowire associated with some SI, the operator must provide the following configuration detail:

- Upstream and Downstream Classifiers
- Upstream and Downstream Service Flows
- NSI L2VPN Encapsulation information - which includes the following:
 - L2VPN Mode (if TRAN-trail is operating in Encapsulation mode)
 - [802.1ad] VID - e.g., TLV 43.5.8.4
 - MPLS Peer
 - Pseudowire ID
 - Service Delimiter

All of the necessary configuration objects may be provisioned via the CM configuration file or there may be a mixture of manual provisioning on the DPoE System and dynamic provisioning via CM configuration file.




Figure 11 - L2VPN Example Config

For this example, assume that all the configuration detail is present within the CM configuration file and assume that the S interface where the SI is to be provisioned is configured to operate as an MU in encapsulation mode.

Figure 11 depicts the basic configuration of the various elements including the D-ONU and sub-elements within the DPoE System. This diagram depicts a simplified view of the configuration output post CM configuration file parsing. In the case where more than one bridge is present in the system it is anticipated that the bridge configuration will have to occur on all internal bridge elements to provide continuity of the TRAN-trail from TUL <> MN_I interface.

In the upstream direction, the provisioned service will classify on the CMIM in the upstream, associate the customer frames with a service flow defining the QoS parameters, and then encapsulate the matching frames with an S-VID of 100. The DPoE System OLT and bridge will forward the service frames matching S-VID=100 as a connection-oriented ethernet bridge from the TUL interface to the MNI at which point the Router will match on the configured Service Delimiter (S-VID=100) and associate those frames to a pseudowire forwarder. In this case we assume the default Pseudowire Class is configured to pop the Service Delimiter. Thus forwarding out the D interface is the original frame as ingressed at the MU interface preceded by the MPLS header and additional layer-2 header.

In the downstream direction, the provisioned service will match on the MPLS label which will include (at a minimum) a VPN label matching the pseudowire instance. The MPLS VPN labels are typically learned dynamically via BGP or LDP. Since the pseudowire forwarder is point-to-point, there is no learning required, the pseudowire forwarder will remove the pseudowire header and push an S-VID equal to the Service Delimiter onto the customer frame before forwarding out the MNI interface. The DPoE System will forward the frame toward the TUL interface as a connection-oriented ethernet bridge and classify on the S-VID in the downstream direction to the correct service-flow and LLID. When the D-ONU receives the frame it will classify on the LLID and S-VID, find the egress MU, then perform the reciprocal decapsulate in the downstream direction to deliver the original customer frame out the correct MU interface.

In this example, the Upstream classifier would match on the CMIM while the Downstream classifier would match on S-VID = 100. The QoS Configuration is similar to that of any MEF service. The additional features come from the L2VPN NSI Encapsulation object. The configuration detail needed to instantiate an MPLS pseudowire is depicted in the example in Figure 12.

L2VPN NSI Encapsulation	Classifiers
43.5.1 VPN Identifier: 10000	US Classifier: CMIM
43.5.2 NSI Encapsulation	DS Classifier: S-VID=100
.3 [802.1ad]=100	SF - QoS Configuration
.4 MPLS Peer=x.x.x.x	US CIR: 100 mbps
43.5.11 Pseudowire ID=xxxxx	DS CIR: 100 mbps
43.5.19 Service Delimiter	US CBS: 32K
.2 S-VID=100	DS CBS: 32K
43.5.18 Pseudowire Class=DEFAULT	Upstream Scheduling: BE

Figure 12 - Example CM Configuration File

8.4.2 [RFC 6074]-Based Service Provisioning

This provisioning model uses [RFC 6074] for VPLS signaling and discovery. With proper configuration [RFC 6074] can be made to provision centralized ELAN forwarding services via VPWS, or distributed forwarding services via VPLS. To support this provisioning model, there are a number of new TLVs required and defined in [DPoE-MEFv2.0], the set of which includes the VSI, VPLS-Class, E-Tree-Role, Root-VID, Leaf-VID, Service Delimiter, and a BGP object with L2VPN-ID Reference, Route Distinguisher, and Import and Export Route Targets.

A brief description of each of the new TLVs is provided here, though the detailed description and interaction between the TLVs is defined in [DPoE-MEFv2.0]. The object-model for the relationship between the existing and new TLVs is shown in Figure 13.




Figure 13 - VSI New TLV Object Model

- VSI: the Virtual Switch Instance is a TLV which includes a number of sub-TLVs. The set of VSI sub-TLVs includes all the values required to provision a new service instance on the DPoE System or modify an existing service instance on the DPoE System to add additional participating TRAN-trails.
- VPLS-Class: the VPLS-Class is used to define standard operator VPLS behavior. The VPLS-Class is a configuration file reference to an object configured on the DPoE System. The VPLS-Class includes objects such as bridge group MTU, service delimiter tag pop or retain actions, and Flow Aware Transport for VPLS, etc. The complete description of the VPLS-Class object is included in [DPoE-IPNEv2.0].
- E-Tree-Role: the E-Tree-Role defines the specific role of the attachment circuit with respect to the bridge group. Available roles are either a root attachment circuit or a leaf attachment circuit.
- Root-VID: On transmission, the Root VID object defines which VLAN ID to apply on frames as they egress towards the D-interface to identify the frame on receipt as originating from a root attachment circuit. On receipt, the Root VID object defines which VID identifies a root-originated frame.
- Leaf-VID: On transmission, the Leaf VID object defines which VLAN ID to apply on frames as they egress towards the D-interface to identify the frame on receipt as originating from a leaf attachment circuit. On receipt, the Leaf VID object defines which VID identifies a leaf-originated frame.
- BGP: The BGP TLV includes a number of sub-options to define the BGP parameters necessary for the service configuration.
- VPN-ID: The reference from the BGP object to the service L2VPN associated with the referenced VPNID object used to identify the BGP object associated with the VPN configuration.
- RD: the Route Distinguisher use within BGP is described by [RFC 4364].
- Route Target (Import): the Import Route Target use within BGP is described by [RFC 4364].
- Route Target (Export): the Export Route Target use within BGP is described by [RFC 4364].

There are three general methods for provisioning and configuring these advanced services. The first is by provisioning the service instance manually, such that an operator might manually configure the VSI on the DPoE System. Based on the configuration file, the DPoE System would then associate the spoke TRAN-trails to the VSI. To provision this service association, the operator would provision the following TLVs in the CM configuration file for the VSI attachment circuit: the VPN-ID the same as the manually-configured VSI, and the Service Delimiter. This would associate the service delimiter to the manually configured VSI. In the case of an E-Tree service, the operator would also provision the E-Tree Role TLV for the given TRAN-trail in addition to the aforementioned TLVs.

The second method for provisioning and configuring these advanced services is by a more complete CM configuration file, which would include the VSI configuration in the CM configuration file. To configure one spoke of an EVPL or ELAN service, the operator would provision the following TLVs in the CM configuration file: VPN-ID, VSI, VPLS-Class, Service Delimiter (SD), and optionally the BGP object to include the Route Distinguisher, RT-Export and RT-Import. For E-Tree services, an additional TLV – the E-Tree-Role TLV – would be configured to provision the spoke as either a LEAF or ROOT. The default would be the same as setting the value to 0 – perform as a root.

The VPN-ID should be a unique, global identifier for the VPLS instance within the operator's network, though there may be multiple spokes on a given DPoE Network that all reference the same VPN-ID. The VPLS-Class would point to a class-name that is pre-configured on the DPoE System. Within the VPLS-Class, the DPoE System would have the definition of the local bridge instance including items like learning behavior (on or off), max MAC address table size, IB-BEB B-DA MAC Translation, and service delimiter behavior (pop or retain).

The DPoE System would use the Service Delimiter TLVs to define the selector-bytes of the [802.1Q] header used by the VSI to associate the TRAN-trail as a logical attachment circuit to the VSI or PW forwarder. The Route Distinguisher (RD) is used to set the BGP extended community value that identifies the VPLS instance within BGP.

E-Tree service construction can be achieved by setting the Leaf and Root VLAN identifiers within the VSI. Further VPLS topological improvements can be made by configuring a separate root and leaf import/export RTs such that only the VSIs that own a root attachment circuit build topological relationships to VSIs that connect to only leaf attachment circuits. In this scenario, VSIs that only connect to leaf attachment circuits will not build attachment circuits to other VSIs that only connect to leaf attachment circuits.

The third method is really a variant, which can be used with either the first or the second method described above. The third method is intended to simplify the CM configuration file by having the DPoE system automatically select the [802.1Q] encapsulation values from the serving group (SG) configured in IPNE which also automatically associates those encapsulation values to the MPLS service element via a service delimiter. The details of this method are further described in Section 8.4.4.

8.4.3 Associating Service Delimiting Tags to an MPLS Service

The router within the DPoE System must have a mechanism for associating a specific TRAN-trail to an attachment circuit associated with a PW forwarder or VSI. The Service Delimiter TLV defines information contained in the [802.1Q] Ethernet frame that the DPoE System uses to identify the frame as belonging to a particular attachment circuit. The DPoE System uses the service delimiter to associate the frame as observed on the MN_I to a provisioned VSI or PW forwarder.

There are two ways to provision a DPoE System to associate a TRAN-trail to a specific VSI or PW forwarder. The first way is by explicitly provisioning the Service Delimiter TLVs which includes the C-VID, S-VID, I-SID, and B-VID. If the operator desired to configure a specific S-VID as a service delimiter on a VSI, the operator would configure only the SD S-VID sub-TLV. Similarly, if an operator desired to configure a specific I-SID and B-VID as a service delimiter to associate it as an attachment circuit on a VSI, the operator would configure the SD I-SID and B-VID sub-TLVs. The SD TLV defines the outer-most [802.1Q] tags as observed at the MN_I by the DPoE System's VSI or Pseudowire Forwarder. The VSI would only match the frame if the sub-TLVs match the outer tags. If an Operator were to provision the SD S-VID sub-TLV and a frame ingressed the VSI at the MN_I with a matching S-VID but also with an I-Tag, the DPoE System VSI would not consider the frame a match. One or more of these TLVs can be provisioned to define the service delimiter.

The second way to provision a DPoE System to associate a TRAN-trail to a specific VSI or PW forwarder is via the simplified MPLS provisioning model where the operator provisions an SG TLV which references an SG identifier configured on the DPoE System, allowing it to allocate the service delimiter (along with the encapsulation values) from the range of values identified in the referenced SG. The SG allows the DPoE System to dynamically allocate the [802.1Q] tags utilized for upstream and downstream encapsulation and forwarding at the D-ONU, and utilized for classification and forwarding at the DPoE System. The configuration of the SG objects is described in [DPoE-IPNEv2.0].

8.4.3.1 Service Delimiter Requirements

The DPoE System MUST support both the provisioning and forwarding behavior to associate a TRAN-trail to a VSI via a provisioned SD S-VID sub-TLV.

The DPoE System MUST support both the provisioning and forwarding behavior to associate a TRAN-trail to a VSI via a provisioned SD I-SID sub-TLV.

The DPoE System MUST support both the provisioning and forwarding behavior to associate a TRAN-trail to a VSI via a provisioned SD B-VID sub-TLV.

The DPoE System MUST support both the provisioning and forwarding behavior to associate a TRAN-trail to a VSI via a provisioned SD I-SID and B-VID sub-TLV.

The DPoE System MUST support forwarding of the frame from MN_i into the VSI or Pseudowire forwarder by matching the outermost two [802.1Q] tags.

The DPoE System MAY support the provisioning and forwarding behavior to associate a TRAN-trail to a VSI via any other set of provisioned SD sub-TLVs.

8.4.4 Simplified MPLS Provisioning Mode for Encapsulation Mode

This provisioning mode functions with Encapsulation Mode only. The simplified mode enables the DPoE System to dynamically select and configure the service delimiter, encapsulation tags, and downstream classifier values to associate the SI with the VSI or PW service instance as defined in the CM configuration file. In the Simplified MPLS Provisioning Mode for Encapsulation Mode, the operator must still provide the TLVs to provision the upstream classifier and the upstream and downstream service flow parameters within the CM configuration file.

In DPoE version 1.0 specifications, as in [L2VPN], the upstream and downstream service flows are not explicitly linked. There has previously been no need to construct a relationship within the CM configuration file between the upstream and downstream service flows as the classification and forwarding behavior is explicitly provisioned in both directions. In the case of the Simplified Provisioning Mode, there is a need to explicitly define the relationship between upstream and downstream service flows to ensure the upstream encapsulation TLVs and downstream classifiers are correctly provisioned as a matched pair. The way the upstream and downstream service flows are associated to each other is described in [DPoE-MULPIv2.0].

8.4.4.1 Simplified PB MPLS Provisioning

For Simplified PB MPLS Provisioning, the mandatory configuration file TLVs include the upstream classifiers, service flow provisioning, VPN-ID, and VSI provisioning objects. The SG TLV (defined in [DPoE-MULPIv2.0]), is used to provide an object the DPoE System can explicitly key on to perform Simplified PB MPLS Provisioning for the service flow.

Within [DPoE-IPNEv2.0] there is an SG configuration item to set aside a range of S-VIDs for use in MPLS auto-provisioning. The Simplified PB MPLS Provisioning mode can be used to support the [L2VPN] auto-provisioning model and the [RFC 6074]-based provisioning model. The following normative statements define the behavior for configuration TLVs specific to each model.

8.4.4.1.1 Simplified PB MPLS Provisioning Requirements

When the vCM parses the CM configuration file and sees an SG TLV in the service flow matching the SG ID of an SG configured on the DPoE System, the DPoE System MUST assign an S-VID from the VPN-SG S-VID pool for use as the encapsulation VID in the upstream and as the downstream classifier for a service flow.

When the vCM parses the configuration file and sees a service flow associated with a VPN-ID, which includes a VSI and a service flow including the SG TLV with a matching ID of an SG configured, the DPoE System MUST assign an S-VID from the SG S-VID pool for use as the encapsulation VID in the upstream and as the downstream classifier for the service flow.

When configured for Simplified PB MPLS Provisioning, the DPoE System MUST use the S-VID selected from the referenced SG S-VID pool for the service flow downstream classifier.

When configured for Simplified PB MPLS Provisioning, the DPoE System MUST use the S-VID selected from the referenced SG S-VID pool for the upstream encapsulation configuration on the D-ONU.

When configured for Simplified PB MPLS Provisioning, the DPoE System MUST use the S-VID selected from the referenced SG S-VID pool for the service delimiter to associate the SI as a local attachment circuit to the VPLS or PW-Forwarder.

8.4.4.2 Simplified PBB MPLS Provisioning for Encap Mode

There is no Simplified PBB MPLS Provisioning model in this version of the specification.

8.5 Local VSI and PW Forwarder Configuration Management Requirements

The DPoE System MUST have the capability of displaying the VSI configured via the CM configuration file in the same way as the DPoE System would display the manually configured VSI.

In the case of the VSI configured via the CM configuration file, the configuration would be volatile and not persist through DPoE System reboots while the manually configured VSI would be stored in non-volatile storage and persist through reboots.

The DPoE System MUST treat the VSI configuration instantiated via CM Configuration file as volatile in that it does not persist through DPoE System reboots.

For instance, if an operator provisioned a D-ONU SI participating in VSI with a VPN-ID of BLUE, when that operator executes a command to show the operating configuration on the DPoE System, the DPoE System MUST display the current running configuration of that VSI.

8.5.1 Adding Attachment Circuits to an Active VSI or PW Instance on a DPoE System

When a new SI is activated and associated with a VSI that already has provisioned attachment circuits, there is a set of behaviors corresponding to each of the two provisioning methods. There is one behavior required when the VSI is manually configured on the DPoE System. There is a different behavior required when the VSI is dynamically configured via the CM configuration file. These normative behaviors are described below.

If, through the dynamic configuration update, a vCM parses a CM configuration file that includes a VPN-ID that is already provisioned on the DPoE System for the vCM and D-ONU undergoing the dynamic configuration update and that VPN-ID provisioned on the DPoE System is associated with an MPLS Peer, the DPoE System MUST replace the configuration of the previously provisioned VPN-ID with the new configuration for the VPN-ID.

If a vCM downloads and parses a configuration file via a new D-ONU registration, and that CM configuration file includes a VPN-ID that is already provisioned on the DPoE System and that VPN-ID provisioned on the DPoE System is associated with an MPLS Peer, the DPoE System MUST reject the configuration file.

If a vCM parses a configuration file that includes a VSI with a VSI-ID that is already provisioned on the DPoE System, the vCM MUST ignore the following TLVs: VPLS-Class, Route-Distinguisher, Root VID, and Leaf VID.

If a vCM parses a configuration file that includes an L2VPN encoding with a VPN-ID that is already provisioned on the DPoE System, the DPoE System MUST act upon the following TLVs by adding them to the already present set on the VSI: Service Delimiter. Note that the service delimiter in this case can be the service delimiter as derived from the Simplified Provisioning mode or one present in the CM configuration file.

As an example of the expected behavior, suppose a D-ONU (D-ONU-A) registers with the DPoE System; the vCM corresponding to D-ONU-A accordingly downloads the configuration file and configures the D-ONU with an MU-SI. That CM configuration file defines that the MU-SI is associated with some VPN-ID, and when the vCM parses the configuration file it discovers the service delimiter that it needs to associate with that VPN-ID has an S-VID=3. Since this is the first time the VSI has been activated on the DPoE system, the vCM configures the DPoE System to activate the VSI with that VPN-ID and associate the VPN-ID with the configured VPLS-Class, and attachment circuit which has an S-VID=3 to the VSI. Subsequently, another D-ONU comes online on the same DPoE System and the vCM, after it downloads the configuration file, finds another MU-SI that is a member of the same VSI (determined by the matching VPN-ID) and has defined an S-VID=27 in the service delimiter TLV. The vCM would configure the VSI to add a new attachment circuit with an S-VID=27 to that VSI.

8.5.2 Modifying an Active VSI or PW Instance on a DPoE System

There are a number of configuration parameters that are not modified via the CM configuration file when the VSI is already instantiated on the DPoE System. These values are: VPLS-Class, Route-Distinguisher, Root VID, and Leaf VID.

In the case that an operator needs to modify these values on a VSI instantiated through the CM configuration process, there are two mechanisms for directly modifying these values. An operator can make changes to the VSI, Bridge Group Instance, or PW instance via SNMP Write Access. Alternatively, an operator can make changes to the VSI, Bridge Group Instance, or PW instance via CLI access.

8.5.3 Removing Attachment Circuits from an Active VSI or PW Instance on a DPoE System

The DPoE System needs to clean up the configuration once certain configuration elements are changed or removed via D-ONU de-registration.

When the D-ONU that has SIs associated with a specific VSI, as denoted by the VPN-ID, de-registers, the vCM MUST remove the attachment circuit association from the VSI.

When the D-ONU that has an SI associated with an MPLS Peer de-registers, the vCM MUST remove the PW configuration associated with that SI.

A DPoE System MUST remove the configuration for a dynamically provisioned VSI that no longer has any member attachment circuits associated with it.

A DPoE System MUST NOT remove the configuration for a manually configured VSI that no longer has any member attachment circuits associated with it.

8.6 L2VPN Layer 2 Forwarding Services

DPoE supports [L2VPN] Layer 2 forwarding mode in point-to-point and multipoint configurations. DPoE expands on the functionality of [L2VPN] multipoint forwarders to allow *or* prohibit forwarding between service flows associated with different D-ONUs (i.e., user-to-user communications). As described in [L2VPN], TLV 43.5.27 is used to define multipoint forwarding mode. The supported [L2VPN] configuration, capabilities, and restrictions are indicated below in Table 5. Sample L2VPN encodings for point-to-point and multipoint configurations are provided in [L2VPN], Appendix I.

Table 5 - L2 Forwarding Modes

Service Type	VPNID: TLV 43.5.1	VLAN ID: TLV22/23.14.2/6	Notes
Multipoint (i.e., simple VLAN)	Same for vCMs	Same for vCMs	MAC learning enabled. User-user communication may be restricted via TLV 43.5.27
Point-to-Point Connection	Same for vCMs	Different among vCMs	No MAC learning. Only 2 connections are permitted. Additional instances are rejected.
Point-to-Point Connection	Different among vCMs	Same for vCMs	No MAC learning. Only one connection using the designated VLAN ID is allowed. Additional instances are rejected.
Point-to-Point Connection	Different among vCMs	Different among vCMs	No MAC learning. Both endpoints may reside within the same DPoE System or a single endpoint can exist outside the DPoE System.

8.7 IP Transport Introduction

To support the construction of EVPL and E-LAN services via VPWS/VPLS, the operator network consists of some interior gateway protocol for local reachability, such as LDP and MP-BGP. The specific requirements of the DPoE System for those protocols are detailed in [DPoE-IPNEv2.0].

8.7.1 General VPWS/VPLS Forwarding

As shown in Figure 14, the DPoE System VPWS/VPLS forwarding system is split into two logical components; 1) the connection-oriented-[802.1Q] bridge module made up of the OLT and PBB/PB forwarder and 2) the VPLS Capable Provider Edge (PE), or VPLS Edge (VE). These elements are connected to each other for Metro Ethernet services via the MN_1 interface.




Figure 14 - VE Reference Model

Within the DPoE specifications, it is a fundamental requirement that the service delimiters are unique within a DPoE Network. There is no requirement of service delimiter uniqueness across multiple DPoE Networks. Specifically, the service-delimiting VIDs (and/or ISIDs), collectively known as the service-delimiters, that identify the TRAN-trail are unique across a set of TULs attached to the DPoE System. For the VE to properly associate a specific TRAN-trail to a service, the VE has to assume uniqueness of the service delimiter on the MN_1 interface. This enables the VE to properly classify on the TRAN-trail via the service delimiters and to associate that TRAN-trail to some VSI or PW forwarder.

The DPoE System and VE MUST use the SD TLV or the auto-allocated service delimiter value (see Section 8.4.4) to identify the frames on the MN_1 interface and associate those frames with the PW forwarder or with the VSI.

If the PW or VSI is configured to pop the service delimiting tag, the DPoE System MUST strip the service delimiters before the frame is sent to the PW per [RFC 4448].

When configured to strip the service delimiters, the DPoE System MUST only strip service delimiters as identified by the SD sub-TLVs.

The DPoE System MUST NOT strip the [802.1Q] I-Tag service delimiter prior to sending the frame to the PW or VSI.

If the "retain tag" configuration value is set within the VPLS-Class, the DPoE System MUST NOT strip the C-Tag, S-Tag, or B-Tag before the frame is sent to the PW.

The DPoE System MUST by default copy the PCP bits from the outer-most tag of the service delimiter into the MPLS Traffic Class field.

Thus, if the service delimiter includes both an I-Tag and a B-Tag, the DPoE System will copy the PCP values from the B-Tag into the MPLS Traffic Class field [RFC 5129].

The DPoE System MUST transpose the PCP values into the MPLS Traffic Class field based on a configuration table as defined in [DPoE-IPNEv2.0].

8.7.2 E-Tree Forwarding

VPLS treats all attachment circuits as equal and provides no distinction between root and leaf. As a result, VPLS forwards in an any-to-any fashion based strictly upon the layer-2 bridge table associated with that VSI. E-Tree adds an additional requirement regarding the ability to restrict certain attachment circuit to attachment circuit forwarding based on port rules. The essential rule for E-Tree forwarding is a restriction of leaf-to-leaf communication. Root-to-root, root-to-leaf, and leaf-to-root forwarding are all permitted in E-Tree.

As a result, with the E-Tree forwarding case, there are a few additional requirements applied to the VSI local to a DPoE System. When a VSI is configured for E-Tree, the VSI is configured to apply either an S-Tag or a B-Tag, depending on the local configuration of the VPLS-Class. In a DPoE System, there is no need to provide both root and leaf identifiers from the SIs, as each SI is distinctly either a root SI or a leaf SI; this is determined when an SI is provisioned and mapped into a VSI.

This is no different than what is required to occur across the provider MPLS network by pushing a specific S-tag with one S-VID to identify a frame that ingress on a root-SI and a different S-VID to identify a frame that ingress on a leaf-SI.




Figure 15 - E-Tree Forwarding

In Figure 15 above, the original frame ingressing into the MU is encapsulated at the D-ONU per provisioning which associates the S-VID as the service delimiter to allow the VE to associate the SI as an attachment circuit of the VSI. The VSI is configured such that root and leaf S-VIDs are defined. The VSI is also informed via provisioning that the attachment circuit is a leaf attachment circuit, and as a result, when the VSI forwards the frame across the provider network, it swaps the service-delimiting S-VID (or B-VID) with the leaf S-VID (or B-VID) and encapsulates the entire Ethernet frame in a pseudowire and MPLS header.

As an example of this forwarding behavior, suppose DPoE System A's VE (shown in Figure 15) receives the frame with the pseudowire header, associates the pseudowire with the proper VSI, notes that the S-VID is a leaf S-VID, and the bridge table shows the C-DA is destined to a root attachment circuit. The VSI strips the PW header, swaps the leaf S-VID with the S-VID encapsulation as defined by the local destination attachment circuit / service

delimiters for the proper MU-SI, and forwards the frame out the MN_i. The D-ONU strips the encapsulation and forwards the original frame out the MU.

In this version of the specification, it is required that an E-Tree service that has multiple VEs participating in the same E-Tree service all use the same S-VID/B-VID to identify root-originated frames and the same S-VID/B-VID to identify leaf-originated frames. The way to ensure that the S-VIDs are the same across a given service is via proper provisioning mechanisms.

Within a local VSI, it is required that forwarding between local attachment circuits is handled by bridging and forwarding logic that restricts leaf-to-leaf communication based on the provisioning of the attachment circuit.

When configured to operate as an E-Tree, the DPoE System VSI MUST NOT forward frames from a leaf attachment circuit to a leaf attachment circuit.

When configured to operate as an E-Tree, the DPoE System VSI MUST NOT forward frames arriving on the D interface with the leaf S-VID (or B-VID) to leaf attachment circuits.

8.7.2.1 PB Learning and Forwarding

There are two learning modes that can occur within a VSI – qualified and unqualified learning. The specific learning mode is operator-provisioned within the VPLS-Class object defined within [DPoE-IPNEv2.0].

The DPoE System MUST support unqualified learning within a VSI.

The DPoE System SHOULD support qualified learning within a VSI.

The DPoE System VSI in PB-forwarding mode MUST perform MAC bridging between the logical attachment circuits (designated by the service delimiter) and the dynamically constructed pseudowires which connect the local VSI to other VSI or PW forwarders across the operator network.

8.7.2.2 PBB Learning and Forwarding

The first PBB forwarding case is one where the DEMARC performs the function of an I-BEB or an IB-BEB, and the D-ONU performs the function of a B-BEB or simply classifies and forwards the frames without participation in PBB. The DPoE System would perform the function of a BCB. This is a simple case that involves basic PB MAC forwarding on the DPoE System as the DEMARC is performing the B-DA learning. This scenario is depicted in Figure 16.




Figure 16 - DPoE-PBB-DEMARC

Referring to Figure 16, note that the Customer Bridge Port (CBP) is connected to the DEMARC in this case and that the DEMARC is operating as the I-BEB. The DPoE System (according to the provisioning of the VSI) can associate

the TRAN-trail to the VSI based on the service delimiter. The service delimiter may identify only the B-VID, only the I-SID, or both I-SID and B-VID, to enable both flexibility and service scaling at the MN_i . Whether the DPoE System is using a service delimiter including only the B-VID, B-VID + I-SID, or only the I-SID, the DPoE System VSI would perform the function of the BCB in this configuration – either MAC forwarding only (unqualified learning) or MAC forwarding + B-VID forwarding (qualified learning).

In the second PBB forwarding case, (see Figure 17), the D-ONU operates as the encapsulation half of an IB-BEB. The D-ONU may be provisioned to add both an I-Component and a B-Component or just an I-Component. The CBP is connected to the D-ONU, which operates as an I-BEB or IB-BEB based on the service configuration. The provisioning needed on the VSI to configure the DPoE System IB-BEB to perform the B-DA translation is within the VPLS-Class, defined in [DPoE-IPNEv2.0].

In the case where the B-DA is not provided in the CM configuration file, the DPoE System will create the Backbone Service Instance Group address and send that to the D-ONU to use as the B-DA for the complete I-Tag encapsulation. As specified in [802.1Q], section 26.3, the Backbone Service Instance Group address is constructed by concatenating the three octets IEEE 802.1Q Backbone Service Instance Group (00-1E-83) with the three octet I-SID, and asserting the I/G bit in the first octet of the resultant value to signify a group MAC address.

The D-ONU adds the encapsulation as provisioned including the Backbone Service Instance Group address as the B-DA and the provisioned I-SID. The DPoE System VSI is then responsible for translating the Backbone Service Instance Group address either to a valid unicast B-DA MAC address (if available) or to the Default Backbone Address as defined by [802.1Q].




Figure 17 - PBB-D-ONU

The DPoE System's VSI MUST support forwarding PBB frames as a BCB.

The DPoE System's VSI SHOULD support I-SID-aware frame forwarding.

As a result of the requirements above, the following requirements are only mandatory when a DPoE System's VSI supports I-SID aware frame forwarding:

- When a configuration file is parsed by a vCM to add an I-Tag encapsulation without a B-DA, the DPoE System MUST create the Backbone Service Instance Group address and provide that to the D-ONU to use as the B-DA for the encapsulation.
- When the DPoE System VSI receives a frame with the Backbone Service Instance Group address, the DPoE System VSI, operating as an IB-BEB, MUST translate the Backbone Service Instance Group address to the actual B-DA if an entry exists for the C-DA within the VSI bridge table.
- If a MAC entry does not exist within the VSI for the Backbone Service Instance Group, the DPoE System MUST translate the B-DA to the Default Backbone Address as defined in [802.1Q].

- The DPoE System VSI MUST NOT forward between two attachment circuits that have different I-SIDs.

If a DPoE System receives a configuration file including an I-SID as a service delimiter and the DPoE System VSI does not support I-SID-aware frame forwarding, the DPoE System MUST reject the configuration file.

8.7.2.2.1 Upstream PBB Forwarding

This section describes the upstream forwarding from the SI toward the D interface including B-DA translation within the VSI.

As shown in Figure 18, a PB-tagged frame arrives at the MU interface and is encapsulated with a PBB encapsulation to include the I-Tag, B-Tag, B-SA, and B-DA. The I-Tag and the B-Tag come from the CM configuration file. The B-SA is manually configured on the DPoE System and sent to the D-ONU via eOAM messages. The B-DA can either be generated by the DPoE System or explicitly provisioned within the CM configuration file. In this forwarding explanation, the B-DA for this service is explicitly left blank and thus the DPoE System generates the Backbone Service Instance Group MAC address and sends that to the D-ONU via eOAM messages. The D-ONU encapsulates the original frame in the provisioned PBB header and sends the frame towards the TU interface.

The DPoE System receives the frame on the TUL interface and forwards across the local PBB bridge based on the B-VID and I-SID to the MN_i. The VE associates attachment circuits based on the provisioned service delimiter for a given VSI – these service delimiters include the B-VID and I-SID. The frame ingresses into the VSI which is provisioned to pop the service delimiter.




Figure 18 - PBB Upstream Forwarding

Based on that VSI configuration, the VSI strips the B-VID, looks at the B-DA which is the Backbone Service Instance Group MAC address, and then performs a lookup in the C-DA to B-DA MAC bridging table for the unicast B-DA that is associated with the C-DA. Assuming the VSI finds the unicast B-DA, the VSI replaces the Backbone Service Instance Group MAC address with the unicast B-DA. In the case that no unicast B-DA is present in the C-DA to B-DA MAC bridging table for the C-DA, the VSI replaces the B-DA of the Backbone Service Instance Group MAC address with the default backbone MAC address.

Once all packet manipulations are performed as described above, the VSI will forward the frame to the egress attachment circuits and/or pseudowires based on the bridge table lookup.

8.7.2.2.2 Downstream PBB Forwarding

This section describes the downstream forwarding from the D-interface to the MU interface including I-SID translation within the VSI. This scenario is shown in Figure 19.




Figure 19 - PBB Downstream Forwarding

In the downstream direction, the MPLS-encapsulated frame arrives at the D interface. This frame may include an LSP label and a VPN label or may only include the VPN label if the upstream router performed Penultimate Hop Popping. The VE associates the ingress frame to a VSI based on the VPN label.

At ingress the VSI strips the PW header and looks up the egress attachment circuit based on the B-DA in the [802.1Q] header. The VSI also pushes the B-Tag onto the frame if the attachment circuit has an associated B-VID and is configured to strip the service delimiter in the upstream direction.

When the frame is transmitted by the VSI and VE toward the MN_i interface in the downstream direction, the frame has the same B-VID and I-SID as used on the local PBB bridge to send the frame toward the correct LLID and ultimately to the correct D-ONU and MU interface.

The DPoE System will forward the frame from across the local DPoE bridge from the MN_i toward the TU interface. The D-ONU will receive the frame and, based on configuration, pop the encapsulation and forward the original frame out the MU interface.

Appendix I Acknowledgements

On behalf of our industry, we would like to thank the following individuals for their contributions to the development of this specification, listed in alphabetical order of company affiliation.

Contributor	Company Affiliation
Mike Emmendorfer	Arris
Dylan Ko	Qualcomm-Atheros
John Dickinson, Edwin Mallette	Bright House Networks
Drew Davis, Paul Gray, Victor Hou	Broadcom
Lane Johnson, Curtis Knittle, Glenn Russell, Vikas Sarawat, Karthik Sundaresan	CableLabs
Jimmy Hu	Ciena
Tim Brophy	Cisco
Jean-Nicolas Huppe	Cogeco
Mehmet Toy, Shamim Akhtar	Comcast
Mike Holmes, Wen Li	Finisar Corporation
James Chen, Hesham ElBakourey	Hitachi
Victor Blake	Independent Consultant
Janet Bean, Nagesh Nandiraju	Motorola
Seth Adler	MRV
Michael Peters, Christopher Griffith	Sumitomo
Robert Harris, Kevin A. Noll, Armin Sepehri	Time Warner Cable
Marek Hajduczenia, Nevin Jones	ZTE

Appendix II Revision History

II.1 Engineering Changes for DPoE-SP-ARCHv2.0-I02-140807

ECN	Date	Summary	Author
ARCHv2.0-N-14.0131-1	3/27/2014	Support for larger MTUs	Marek Hajduczenia
ARCHv2.0-N-14.0140-1	4/11/2014	IP(HSD) changes for Arch2.0	Steve Burroughs
ARCHv2.0-N-14.0169-1	7/03/2014	Alignment and cleanup of 802.3 references	Marek Hajduczenia
ARCHv2.0-N-14.0185-1	7/10/2014	DPoEv2 ARCH edits to Support 2G-EPON	Lane Johnson

II.2 Engineering Change for DPoE-SP-ARCHv2.0-I03-150319

ECN	Date	Summary	Author
ARCHv2.0-N-14.0196-1	10/30/2014	Support for Power Saving mechanism	Marek Hajduczenia

II.3 Engineering Changes for DPoE-SP-ARCHv2.0-I04-150611

ECN	Date	Summary	Author
ARCHv2.0-N-15.0206-1	4/02/2015	Support for Power Saving mechanism	Arkin Aydin
ARCHv2.0-N-15.0209-1	4/30/2015	Enhance L2VPN Multipoint Forward	Arkin Aydin

II.4 Engineering Changes for DPoE-SP-ARCHv2.0-I05-160602

ECN	Date	Summary	Author
ARCHv2.0-N-15.0222-1	12/10/2015	Support for Tree / Trunk optical line protection mechanism	Arkin Aydin
ARCHv2.0-N-15.0225-1	12/31/2015	Remove DEMARC Specification References and Attributes	Arkin Aydin
ARCHv2.0-N-16.0235-1	3/31/2016	DPoE 2.0 ARCH - Retire SOAM Specification	Steve Burroughs

II.5 Engineering Change for DPoE-SP-ARCHv2.0-I06-180228

ECN	Date	Summary	Author
ARCHv2.0-N-18.0259-1	2/8/2018	S/S1/S2 interface alignment and simplification	Marek Hajduczenia

II.6 Engineering Change for DPoE-SP-ARCHv2.0-I07-190213

ECN	Date	Summary	Author
ARCHv2.0-N-18.0268-2	11/21/2018	Add DPoE Embedded Test Agent to Architecture diagram	Steve Burroughs

II.7 Engineering Change for DPoE-SP-ARCHv2.0-I08-230322

ECN	Date	Summary	Author
ARCHv2.0-N- 23.0278-1	3/2/2023	Remove reference to 802.1d and replace with 802.1Q as needed	Steve Burroughs

* * *