

L2VPN WG
Internet Draft
Intended status: Informational
Expires: November 2014

Raymond Key (editor)
Lucy Yong, Huawei (editor)
Simon Delord
Telstra
Frederic Jounay, Orange CH
Lizhong Jin
May 21, 2014

**A Framework for Ethernet Tree (E-Tree) Service over a Multiprotocol
Label Switching (MPLS) Network
draft-ietf-l2vpn-etree-frwk-05.txt**

Abstract

This document describes an Ethernet-Tree (E-Tree) solution framework for supporting the Metro Ethernet Forum (MEF) E-Tree service over a Multiprotocol Label Switching (MPLS) network. The objective is to provide a simple and effective approach to emulate E-Tree services in addition to Ethernet LAN (E-LAN) services on an existing MPLS network.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on November 21, 2014.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction.....	3
1.1.	Conventions used in this document.....	3
1.2.	Terminology.....	3
2.	Overview.....	4
2.1.	Ethernet Bridge Network.....	4
2.2.	MEF Multipoint Ethernet Services: E-LAN and E-Tree.....	4
2.3.	IETF L2VPN.....	5
2.3.1.	Virtual Private LAN Service (VPLS).....	5
2.3.2.	Ethernet VPN (EVPN).....	5
2.3.3.	Virtual Private Multicast Service (VPMS).....	6
3.	E-Tree Architecture Reference Model.....	6
4.	E-Tree Use Cases.....	8
5.	L2VPN Gaps for Emulating MEF E-Tree Service.....	10
5.1.	No Differentiation on AC Role.....	10
5.2.	No AC Role Indication or Advertisement.....	10
6.	Security Considerations.....	11
7.	IANA Considerations.....	11
8.	References.....	11
8.1.	Normative References.....	11
8.2.	Informative References.....	12
9.	Acknowledgments.....	12

1. Introduction

This document describes an Ethernet-Tree (E-Tree) solution framework for supporting the Metro Ethernet Forum (MEF) E-Tree service over a Multiprotocol Label Switching (MPLS) network. The objective is to provide a simple and effective approach to emulate E-Tree services in addition to Ethernet LAN (E-LAN) services on an existing MPLS network.

This document extends the existing IETF specified Layer 2 Virtual Private Network (L2VPN) framework [[RFC4664](#)] to provide the emulation of E-Tree services over an MPLS network. It specifies the E-Tree architecture reference model and describes the corresponding functional components. It also points out the gaps and required extension areas in existing L2VPN solutions such as Virtual Private LAN Service (VPLS)[[RFC4761](#)][[RFC4762](#)] and Ethernet Virtual Private Network (EVPN)[[EVPN](#)] for supporting E-Tree services.

1.1. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119](#) [[RFC2119](#)].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying [RFC-2119](#) significance.

1.2. Terminology

This document adopts all the terminologies defined in [RFC4664](#) [[RFC4664](#)], [RFC4761](#) [[RFC4761](#)], and [RFC4762](#) [[RFC4762](#)]. It also uses the following terminologies:

Leaf Attachment Circuit (AC): An AC with Leaf role. An ingress Ethernet frame at a Leaf AC (Ethernet frame arriving over an AC at the provider edge (PE) of an MPLS network) can only be delivered to one or more Root ACs in an E-Tree service instance. An ingress Ethernet frame at a Leaf AC MUST NOT be delivered to any Leaf ACs in the E-Tree service instance.

Root AC: An AC with Root role. An ingress Ethernet frame at a Root AC can be delivered to one or more of the other ACs in the associated E-Tree service instance.

E-Tree: An Ethernet VPN service in which each AC is assigned the role of a Root or Leaf. The forwarding rules in E-Tree are: Root AC can

communicate with other Root ACs and Leaf ACs; Leaf ACs can only communicate with Root ACs.

2. Overview

2.1. Ethernet Bridge Network

In this document, Ethernet bridge network refers to the Ethernet bridge/switch network defined in IEEE802.1Q [[IEEE802.1Q](#)]. In a bridge network, a data frame is an Ethernet frame; data forwarding is based on destination MAC address; MAC reachability is learned in the data plane based on the source MAC address and the port (or tagged port) on which the frame arrives; and the MAC aging mechanism is used to remove inactive MAC addresses from the MAC forwarding table on an Ethernet switch.

Data frames arriving at a switch may be destined to known unicast MAC destinations, unknown, multicast, or broadcast MAC destinations. Unknown, multicast, and broadcast frames are forwarded in a similar way, i.e. to every port except the ingress port on which the frame arrives. Multicast forwarding can be further constrained when using multicast control protocol snooping or multicast MAC registration protocols. [[IEEE802.1Q](#)]

An Ethernet host receiving an Ethernet frame checks the destination address in the frame to decide whether it is the intended destination.

2.2. MEF Multipoint Ethernet Services: E-LAN and E-Tree

MEF6.2 [[MEF6.2](#)] defines two multipoint Ethernet Service types:

- o E-LAN (Ethernet LAN), a multipoint-to-multipoint service
- o E-Tree (Ethernet Tree), a rooted-multipoint service

According to MEF's technical specification, a multipoint Ethernet service is always bidirectional, which means that any AC in the service can send and receive Ethernet frames to/from customer equipment (CE). Note that the term AC is equivalent to MEF User-Network Interface (UNI). Furthermore, MEF also defines AC roles. One role is Root and another is Leaf. Besides destination MAC-based forwarding, additional forwarding rules defined by MEF for a multipoint Ethernet Service are below:

- o A Root AC can receive/transmit a frame from/to any other ACs.

- o A Leaf AC can receive/transmit a frame from/to any Root ACs.
- o A Leaf AC cannot receive/transmit a frame from/to any Leaf ACs.

For an E-LAN service, all ACs have the Root role, which means that any AC can communicate with other ACs in the service. The E-LAN service defined by MEF may be implemented by IETF L2VPN solutions such as VPLS and EVPN [[EVPN](#)].

An E-Tree service has one or more Root ACs and many Leaf ACs. An E-Tree service supports the communication among the roots and between a root and a leaf but prohibits the communication among the leaves. Existing IETF L2VPN solutions can't support the E-Tree service. This document specifies the E-Tree architecture reference model that supports the E-Tree service defined by MEF [[MEF6.2](#)]. [Section 4](#) will discuss different E-Tree use cases.

[2.3.](#) IETF L2VPN

[2.3.1.](#) Virtual Private LAN Service (VPLS)

VPLS [[RFC4761](#)] [[RFC4762](#)] is an L2VPN solution that provides multipoint-to-multipoint Ethernet connectivity across IP/MPLS networks. VPLS emulates traditional Ethernet Virtual LAN Services (VLAN) in MPLS networks, and may support MEF E-LAN services.

A data frame in VPLS is an Ethernet frame. Data forwarding in a VPLS instance is based on the destination MAC address and the VLAN on which the frame arrives. MAC reachability learning is performed in the data plane based on the source address and the AC or Pseudowire (PW) on which the frame arrives. MAC aging is also the mechanism used to remove inactive MAC addresses from a VPLS switching instance (VSI) on a Provider Edge (PE). VPLS supports forwarding for known unicast, unknown unicast, broadcast, and multicast Ethernet frames.

Many service providers have deployed VPLS in their networks to provide L2VPN services to customers.

[2.3.2.](#) Ethernet VPN (EVPN)

Ethernet VPN [[EVPN](#)] is an enhanced L2VPN solution that emulates an Ethernet LAN or virtual LAN(s) across MPLS networks.

EVPN supports active-active multi-homing of CEs and uses Multiprotocol Border Gateway Protocol (MP-BGP) control plane to advertise MAC address reachability from an ingress PE to egress PEs. Thus, a PE learns MAC addresses reachable over local ACs in the data

plane and other MAC addresses reachable across the MPLS network over remote ACs via the EVPN MP-BGP control plane. As a result, EVPN aims to support large-scale L2VPN with better resiliency compared to VPLS.

EVPN is relatively new technique and is still under development in IETF L2VPN WG.

2.3.3. Virtual Private Multicast Service (VPMS)

VPMS [[VPMS](#)] is an L2VPN solution that provides point-to-multipoint connectivity across MPLS networks and supports various attachment circuit (AC) types, including Frame Relay, ATM, Ethernet, PPP, etc.

In the case of Ethernet ACs, VPMS provides single coverage of receiver membership, i.e. there is no distinct differentiation for multicast groups in one VPN. Destination address in the Ethernet frame is not used in data forwarding.

VPMS supports unidirectional point-to-multipoint transport from a sender to multiple receivers and MAY support reverse transport in a point-to-point manner.

3. E-Tree Architecture Reference Model

Figure 1 illustrates the E-Tree architecture reference model. Three provider edges (PEs), PE1, PE2 and PE3 are shown in the figure. Each of these PEs has a Virtual Service Instance (VSI) associated with an E-Tree service instance. A CE attaches to a VSI on a PE via an AC. Each AC MUST be configured as a root or leaf AC. In figure 1, AC1 AC2, AC5, AC6, AC9, AC10 are Root ACs; AC3, AC4, AC7, AC8, AC11, AC12 are Leaf ACs. This implies that an Ethernet frame from CE01, CE02, etc will be treated as a frame originated from a Root AC; a frame from CE03, CE04, etc will be treated as a frame originated from a Leaf AC.

Under this architecture model, the forwarding rules among ACs, regardless whether sending AC and receiving AC are on the same PE or on different PEs, are described as follow:

- o An egress frame (frame to be transmitted over an AC) at an AC with Root role MUST be the result of an ingress frame at an AC (frame received at an AC) that has Root or Leaf role attached to the same E-tree service instance.
- o An egress frame at the AC with Leaf role MUST be the result of an ingress frame at an AC that has Root role attached to the same E-tree service instance.

These rules apply to all frame types, i.e. Known Unicast, Unknown, Broadcast, and Multicast. For Known Unicast frames, forwarding in a VSI context is based on the destination MAC address.

A VSI on a PE corresponds to an E-Tree instance and maintains a MAC forwarding table which is isolated from other VSI tables on the PE. It also keeps the track of local AC roles. The VSI receives a frame from an AC or across the MPLS core; and forwards the frame to another AC over which the destination is reachable according to the VSI forwarding table and forwarding rules described above. When the target AC is on a remote PE, the VSI forwards the frame to the remote PE over the MPLS core. Forwarding over the MPLS core will be dependent on the E-tree solution. For instance, a solution may adopt PWs to mesh VSIs as in VPLS, and forward frames over VSIs subject to the E-tree forwarding rules. Alternatively, a solution may adopt the EVPN forwarding paradigm constrained by the E-tree forwarding rules. Thus, solutions that satisfy the E-tree requirements could be extensions to VPLS and EVPN.

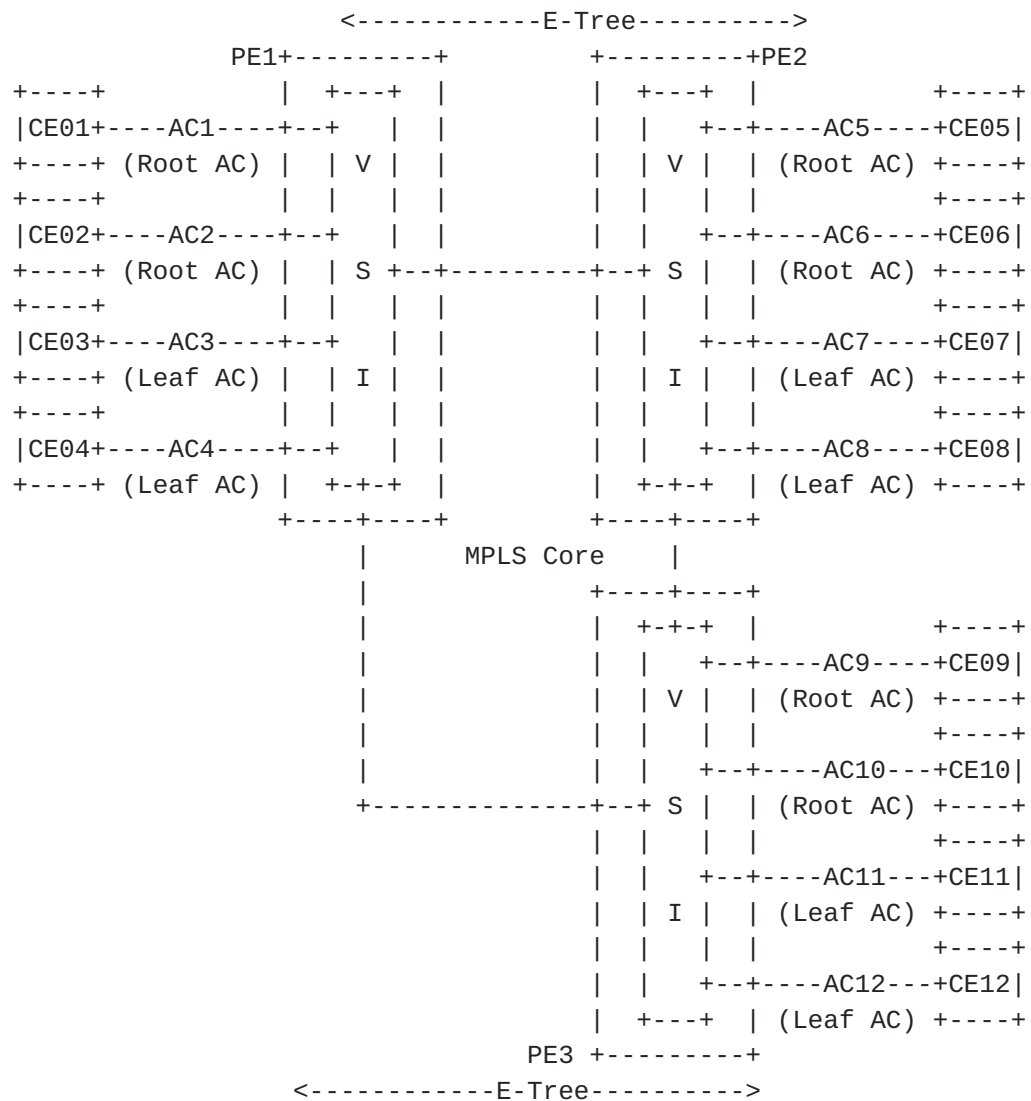


Figure 1 E-Tree Architecture Reference Model

In most use cases, an E-Tree service has only a few Root ACs (root CE sites) but many Leaf ACs (leaf CE sites). Furthermore, a PE may have only Root ACs or only Leaf ACs. Figure 1 provides a general E-Tree architecture model.

4. E-Tree Use Cases

Table 1 below presents some major use cases for E-Tree.

	Use Case	Root AC	Leaf AC
1	Hub & Spoke VPN	Hub Site	Spoke Site
2	Wholesale Access	Customer's Interconnect	Customer's Subscriber
3	Mobile Backhaul	RAN NC	RAN BS
4	IEEE 1588 PTPv2 [1588] Clock Synchronization	PTP Server	PTP Client
5	Internet Access Reference: [TR-101]	BNG Router	Subscriber
6	Broadcast Video (unidirectional only)	Video Source	Subscriber
7	Broadcast/Multicast Video plus Control Channel	Video Source	Subscriber
8	Device Management	Management System	Managed Device

Where:

RAN: Radio Access Network

NC: Network Controller

BS: Base Station

PTP: Precision Time Protocol

BNG: Broadband Network Gateway

Table 1 E-Tree Use Cases

Common to all use cases, direct Layer2 Leaf-to-Leaf communication is required to be prohibited. For Mobile backhaul, this may not be valid for LTE X2 interfaces in the future. If direct Layer 2 Leaf-to-Leaf communication needs to be prohibited, E-Tree should be used.

Also common to the use cases mentioned above, there may be single or multiple Root ACs in one E-Tree service. The need of multiple Root-ACs may be driven by redundancy requirement or multiple serving sites. Whether a particular E-Tree service needs to support single or multiple Root ACs depends on an application.

An E-Tree service can meet these use case requirements.

Among the use cases mentioned above, broadcast video draws most attention. In fact, broadcast video represents an example for broadcast/multicast content delivery in general, such as news feed, financial data feed, etc.

5. L2VPN Gaps for Emulating MEF E-Tree Service

E-Tree Service defines special forwarding rules that prohibit forwarding Ethernet frames among leaves. This poses some challenges to IETF L2VPN solutions such as VPLS and EVPN in emulating E-Tree service over MPLS networks. There are two major issues described in the following sections.

5.1. No Differentiation on AC Role

IP/MPLS L2VPN architecture only has single-role Attachment Circuit (AC) and supports any-to-any connectivity among all ACs. It does not include any mechanism for any forwarding constraint based on the AC role. However, E-Tree service defines two AC roles, Root and Leaf, and defines the forwarding rules among the ACs based on the AC roles.

5.2. No AC Role Indication or Advertisement

In IETF L2VPN, when a PE, say PE2, receives a frame from another PE, say PE1, across the MPLS core, PE2 does not know if the frame is originated from a root AC or leaf AC on PE1. This causes a forwarding issue on PE2 because E-Tree forwarding rules require that the forwarder MUST know the role of the frame origin, i.e. from root AC or leaf AC. This is specifically important, when PE2 has both a root AC and a leaf AC attached to the same VSI. The forwarding rules apply to all types of frames (known unicast destination, unknown unicast destination, multicast and broadcast).

5.3. Other issues

In order to fulfill E-Tree service definition, L2VPN extension is required to support Root and Leaf ACs and communicate AC role among the PEs in an E-Tree instance. Furthermore, some desirable requirements for IETF E-Tree are specific to an IP/MPLS L2VPN implementation such as Leaf-only PE. A Leaf-only PE is the PE that has all ACs associating to a VSI as the Leaf role, thus other PEs on the same E-Tree instance do not necessarily forward the frames coming from Leaf ACs to the Leaf-only PE, which may save some network resources. It is also desirable for E-Tree solution to work with existing PEs, i.e. PEs that only have single-role ACs and the role is equivalent to the root role in an E-Tree Service. These requirements are described in the E-Tree requirement document. [[RFC7152](#)]

6. Security Considerations

There could be cases where an E-tree service is deployed for security reasons to prohibit communication among sites (leaves) and only allow communication between branch sites (leaves) and hub sites (roots). Thus, an E-tree solution must enable the enforcement of an E-tree service definition and the corresponding forwarding constraints. The E-Tree solution must also guarantee that Ethernet frames do not leak outside of the E-tree service instance to which they belong. If additional security is desired, the PE-to-PE tunnels can be IPsec tunnels. For more security, the end systems at the CE sites may use appropriate means of encryption to secure their data even before it enters the Service Provider network.

7. IANA Considerations

The document requires no IANA action.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [IEEE802.1Q] IEEE802.1, "Media Access Control (MAC) Bridges and Virtual Bridged Local Area", IEEE802.1Q, 2011
- [1588] IEEE 1588, "Precision Time Protocol", IEEE 1588, 2013
- [MEF6.2] MEF, "Metro Ethernet Forum, Ethernet Services Definitions - Phase 2", April 2008
- [RFC4664] Andersson, L., et al, "Framework for Layer 2 Virtual Private Network (L2VPNs)", [RFC4664](#), Sept. 2006
- [RFC4761] Kompella & Rekhter, "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC4761](#), January 2007
- [RFC4762] Lasserre & Kompella, "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC4762](#), January 2007
- [RFC7152] Key, et al., "Requirements for Metro Ethernet Forum (MEF) Ethernet-Tree (E-Tree) Support in L2VPN", [RFC7152](#), April 2011.

8.2. Informative References

- [TR-101] Broadband Forum, "Migration to Ethernet-Based Broadband Aggregation Issue 2", July 2011
- [VPMS] Kamite, et al., "Framework and Requirements for Virtual Private Multicast Service (VPMS)", [draft-ietf-l2vpn-vpms-frmwk-requirements-05](#), work in progress
- [EVPN] Sajassi, et al., "BGP MPLS Based Ethernet VPN", [draft-ietf-l2vpn-evpn-07](#), work in progress

9. Contributing Authors

The following people contribute the document as co-authors.

Yuji Kamite
NTT Communications Corporation
Granpark Tower
3-4-1 Shibaura, Minato-ku
Tokyo 108-8118, Japan
Email: y.kamite@ntt.com

Wim Henderickx
Alcatel-Lucent
Copernicuslaan 50
2018 Antwerp, Belgium
Email: wim.henderickx@alcatel-lucent.com

10. Acknowledgments

Authors like to thank Nabil Bitar for this detail review and suggestions.

Authors' Addresses

Raymond Key (editor)

Email: raymond.key@ieee.org

Lucy Yong (editor)
Huawei USA

Email: lucy.yong@huawei.com

Simon Delord
Telstra

Email: simon.delord@gmail.com

Frederic Jounay
Orange CH
4 rue caudray 1020 Renens
Switzerland

Email: frederic.jounay@orange.ch

Lizhong Jin

Email: lizho.jin@gmail.com