

Internet Engineering Task Force (IETF)
Request for Comments: 7387
Category: Informational
ISSN: 2070-1721

R. Key, Ed.

L. Yong, Ed.
Huawei
S. Delord
Telstra
F. Jounay
Orange CH
L. Jin
October 2014

A Framework for Ethernet Tree (E-Tree) Service
over a Multiprotocol Label Switching (MPLS) Network

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 document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7387>.

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. 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	9
5.1. No Differentiation on AC Role	9
5.2. No AC Role Indication or Advertisement	9
5.3. Other Issues	9
6. Security Considerations	10
7. References	11
7.1. Normative References	11
7.2. Informative References	11
Acknowledgments	12
Contributors	12
Authors' Addresses	13

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. Terminology

This document adopts all the terminologies defined in RFC 4664 [RFC4664], RFC 4761 [RFC4761], and RFC 4762 [RFC4762]. It also uses the following terms:

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 an E-Tree are as follows:

- o The Root AC can communicate with other Root ACs and Leaf ACs.
- o 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 IEEE 802.1Q [IEEE802.1Q]. In a bridge network, a data frame is an Ethernet frame; data forwarding is based on destination Media Access Control (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, unknown unicast, multicast, or broadcast MAC destinations. Unknown unicast, 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 using 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

MEF 6.1 [MEF6.1] 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

The MEF defines User-Network Interface (UNI) in a multipoint service as the Ethernet interface between Customer Equipment (CE) and a Provider Edge (PE), i.e., the PE can send and receive Ethernet frames to/from the CE. The MEF also defines UNI roles in a multipoint service. One role is Root, and another is Leaf.

Note that MEF UNI in a service is equivalent to the Attachment Circuit (AC) defined in L2VPN [RFC4664]. The Root AC and Leaf AC defined in this document are the same as the Root UNI and Leaf UNI as defined in MEF 10.3 [MEF10.3]. The terms "Root AC" and "Leaf AC" are used in the following MEF service description.

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 the 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 at least two Leaf ACs. An E-Tree service supports communication among the roots and between a root and a leaf but prohibits 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 the MEF [MEF6.1]. 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 (VLAN) services 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 the mechanism used to remove inactive MAC addresses from a VPLS switching instance (VSI) on a PE. VPLS supports forwarding for known unicast frames, as well as 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 multihoming of CEs and uses the 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 that are 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 a relatively new technique and is still under development in the 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 differentiation among multicast groups in one VPN. The destination address in the Ethernet frame is not used in data forwarding.

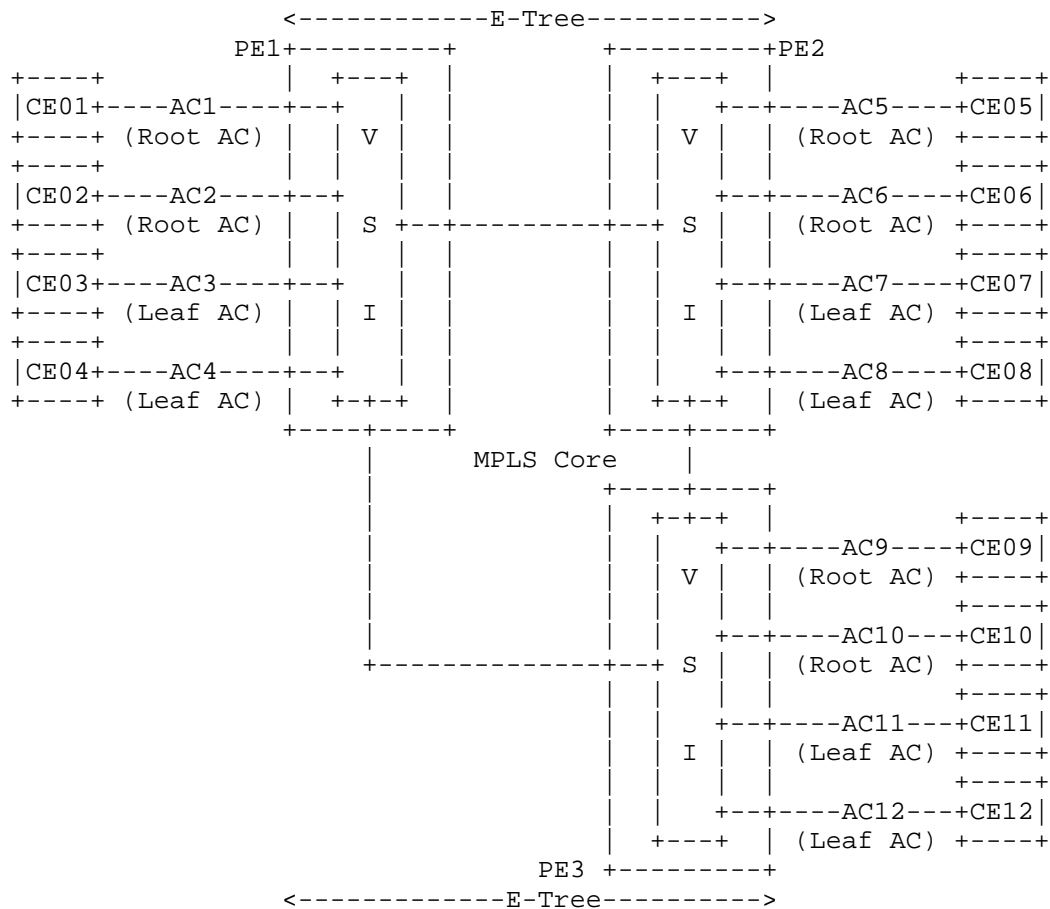
VPMS supports unidirectional point-to-multipoint transport from a sender to multiple receivers and may support reverse transport in a point-to-point manner.

3. E-Tree Architecture Reference Model

Figure 1 illustrates the E-Tree architecture reference model. Three Provider Edges -- PE1, PE2, and PE3 -- are shown in the figure. Each PE has a Virtual Service Instance (VSI) associated with an E-Tree service instance. A CE attaches to the VSI on a PE via an AC. Each AC must be configured with a Root or Leaf role. In Figure 1, AC1, AC2, AC5, AC6, AC9, and AC10 are Root ACs; AC3, AC4, AC7, AC8, AC11, and AC12 are Leaf ACs. This implies that a PE (local or remote) processes the Ethernet frames from CE01, CE02, etc., as if they originated from a Root AC, and it processes the Ethernet frames from CE03, CE04, etc., as if they originated from a Leaf AC.

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

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



These rules apply to all frame types, i.e., known unicast, unknown unicast, 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 service instance and maintains a MAC forwarding table that is isolated from other VSI tables on the PE. It also keeps track of local AC roles. The VSI receives a frame from an AC or across the MPLS core, and it 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 to 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.

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 [IEEE1588] 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 Layer 2 leaf-to-leaf communication is required to be prohibited. For mobile backhaul, this may not be valid for Long Term Evolution (LTE) X2 interfaces; an LTE X2 interface [LTE] enables communication between two evolved Node Bs (eNBs). E-Tree service is appropriate for such use cases.

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

5. L2VPN Gaps for Emulating MEF E-Tree Service

The MEF 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 an MPLS network. There are two major issues described in the following subsections.

5.1. No Differentiation on AC Role

IP/MPLS L2VPN architecture has no distinct roles on Attachment Circuits (ACs) and supports any-to-any connectivity among all ACs. It does not have any mechanism to support forwarding constraints based on an AC role. However, the MEF E-Tree service defines two AC roles -- Root and Leaf -- and defines the forwarding rules based on the originating and receiving AC roles of a given frame.

5.2. No AC Role Indication or Advertisement

In an L2VPN, when a PE, say PE2, receives a frame from another PE, say PE1, over the MPLS core, PE2 does not know if the frame from PE1 is originated from a Root AC or Leaf AC. This causes the forwarding issue on PE2 because the 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 Root AC and Leaf AC attached to the VSI. E-Tree forwarding rules apply to all types of frames (known unicast destination, unknown unicast destination, multicast, and broadcast).

5.3. Other Issues

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 a PE that only has Leaf AC(s) in an E-Tree service instance; thus, other PEs on the same E-Tree service instance do not necessarily forward the frames originated from a Leaf AC to the Leaf-only PE,

which may save some network resources. It is also desirable for an E-Tree solution to work with existing PEs that support single-role ACs, where the role is equivalent to the root in an E-Tree service. These requirements are described in the E-Tree requirement document [RFC7152].

6. Security Considerations

An E-Tree service may be deployed for security reasons to prohibit communication among sites (leaves). An E-Tree solution must enforce E-Tree forwarding constraints. The solution must also guarantee that Ethernet frames do not leak outside of the E-Tree service instance to which they belong.

An E-Tree service prohibits communication among leaf sites but does not have knowledge of higher-layer security constraints. Therefore, in general, higher-layer applications cannot rely on E-Tree to provide security protection unless all security constraints are fully implemented by the E-Tree service.

Enhancing L2VPN for E-Tree services inherits the same security issues described in the L2VPN framework document [RFC4664]. These relate to both control-plane and data-plane security issues that may arise in the following areas:

- o issues fully contained in the provider network
- o issues fully contained in the customer network
- o issues in the customer-provider interface network

The framework document has substantial discussions on the security issues and potential solutions to address them. An E-Tree solution must consider these issues and address them properly. VPLS [RFC4761] [RFC4762] and/or EVPN [EVPN] will likely be candidate solutions for an E-Tree service over an MPLS network. The security capabilities built into those solutions will be naturally adopted when supporting E-Tree. For details, see the Security Considerations sections in [RFC4761], [RFC4762], and [EVPN].

7. References

7.1. Normative References

- [MEF6.1] Metro Ethernet Forum, "Ethernet Services Definitions - Phase 2", MEF 6.1, April 2008.
- [MEF10.3] Metro Ethernet Forum, "Ethernet Service Attributes - Phase 3", MEF 10.3, October 2013.
- [RFC4664] Andersson, L., Ed., and E. Rosen, Ed., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, September 2006, <<http://www.rfc-editor.org/info/rfc4664>>.
- [RFC4761] Kompella, K., Ed., and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007, <<http://www.rfc-editor.org/info/rfc4761>>.
- [RFC4762] Lasserre, M., Ed., and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, January 2007, <<http://www.rfc-editor.org/info/rfc4762>>.
- [RFC7152] Key, R., Ed., DeLord, S., Jounay, F., Huang, L., Liu, Z., and M. Paul, "Requirements for Metro Ethernet Forum (MEF) Ethernet-Tree (E-Tree) Support in Layer 2 Virtual Private Network (L2VPN)", RFC 7152, March 2014, <<http://www.rfc-editor.org/info/rfc7152>>.

7.2. Informative References

- [IEEE802.1Q] IEEE, "IEEE Standard for Local and metropolitan area networks -- Media Access Control (MAC) Bridges and Virtual Bridged Local Area", IEEE Std 802.1Q, 2011.
- [IEEE1588] IEEE, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", IEEE Std 1588, July 2008.
- [LTE] 3GPP, "Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN)", 3GPP TS 36.300 v11.2.0, July 2012.
- [TR-101] Broadband Forum, "Migration to Ethernet-Based Broadband Aggregation", TR-101 Issue 2, July 2011.

- [VPMS] Kamite, Y., Jounay, F., Niven-Jenkins, B., Brungard, D., and L. Jin, "Framework and Requirements for Virtual Private Multicast Service (VPMS)", Work in Progress, draft-ietf-l2vpn-vpms-frmwk-requirements-05, October 2012.
- [EVPN] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., and J. Uttaro, "BGP MPLS Based Ethernet VPN", Work in Progress, draft-ietf-l2vpn-evpn-11, October 2014.

Acknowledgments

The authors would like to thank Nabil Bitar and Adrian Farrel for their detailed review and suggestions.

Contributors

The following people contributed significantly to this document.

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

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

