

L2VPN Workgroup
INTERNET-DRAFT
Intended Status: Standards Track

Ali Sajassi
Samer Salam
Cisco

Expires: December 29, 2012

June 29, 2012

E-TREE Support in E-VPN
draft-sajassi-l2vpn-evpn-etree-00

Abstract

The Metro Ethernet Forum (MEF) has defined a rooted-multipoint Ethernet service known as Ethernet Tree (E-Tree). [ETREE-FRAMEWORK] proposes a solution framework for supporting this service in MPLS networks. This document discusses how those functional requirements can be easily met with E-VPN.

Status of this Memo

This Internet-Draft is submitted to IETF 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/lid-abstracts.html>

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

Copyright and License Notice

Copyright (c) 2012 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	E-Tree Scenarios and E-VPN Support	3
2.1	Scenario 1: Leaf OR Root site(s) per PE	3
2.2	Scenario 2: Leaf AND Root site(s) per PE	5
2.3	Scenario 3: Leaf AND Root site(s) per Ethernet Segment	6
3	Operation	7
3.1	E-Tree with MAC Learning	7
3.2	E-Tree without MAC Learning	8
4	Acknowledgement	8
5	Security Considerations	8
6	IANA Considerations	8
7	References	8
7.1	Normative References	8
7.2	Informative References	8
	Authors' Addresses	9

1 Introduction

The Metro Ethernet Forum (MEF) has defined a rooted-multipoint Ethernet service known as Ethernet Tree (E-Tree). In an E-Tree service, endpoints are labeled as either Root or Leaf sites. Root sites can communicate with all other sites. Leaf sites can communicate with Root sites but not with other Leaf sites.

[ETREE-FRAMEWORK] proposes the solution framework for supporting E-Tree service in MPLS networks. The document identifies the functional components of the overall solution to emulate E-Tree services in addition to Ethernet LAN (E-LAN) services on an existing MPLS network.

[EVPN] is a solution for multipoint L2VPN services, with advanced multi-homing capabilities, using BGP for distributing customer/client MAC address reach-ability information over the MPLS/IP network.

This document discusses how the functional requirements for E-Tree service can be easily met with E-VPN.

1.1 Terminology

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].

2 E-Tree Scenarios and E-VPN Support

In this section, we will categorize support for E-Tree into three different scenarios, depending on the nature of the site association (Root/Leaf) per PE or per Ethernet Segment:

- Leaf OR Root site(s) per PE
- Leaf AND Root site(s) per PE
- Leaf AND Root site(s) per Ethernet Segment

For each scenario, we will describe the E-VPN mechanism for supporting the E-Tree service.

2.1 Scenario 1: Leaf OR Root site(s) per PE

In this scenario, a PE may have Root sites OR Leaf sites for a given VPN instance, but not both concurrently. The PE may have both Root and Leaf sites albeit for different VPNs. Every Ethernet Segment

Value = 6 bytes uniquely identifying an EVI.

This extended community is a new transitive extended community, and will be referred to as the EVI-Import Extended Community. This extended community is used in lieu of the RT on the following E-VPN routes:

- MAC Advertisement Routes
- Ethernet A-D Routes
- Inclusive Multicast Routes

On a per EVI basis, every PE exports routes with the single EVI-Import extended community associated with its type of site(s). Furthermore, a PE with Root site(s) imports routes with both Root and Leaf EVI-Import extended community. Whereas, a PE with Leaf site(s) only imports the Root EVI-Import extended community.

2.2 Scenario 2: Leaf AND Root site(s) per PE

In this scenario, a PE may have a set of one or more Root sites AND a set of one or more Leaf sites for a given VPN instance. Every Ethernet Segment connected to the PE is uniquely identified as either a Root or a Leaf site.

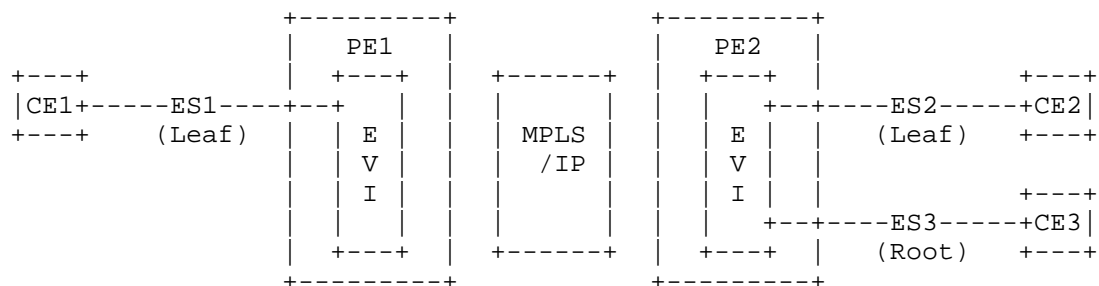


Figure 2: Scenario 2

This scenario requires that the MPLS-encapsulated frames be tagged with an indication of whether they originated from a Root or a Leaf Ethernet Segment, so that the proper connectivity constraints can be enforced. This can be achieved in E-VPN through the use of the ESI MPLS label, since this label identifies the Ethernet Segment of origin of a given frame. For E-Tree service, the ESI MPLS label must be used to encapsulate not only multi-destination frames (i.e. broadcast, multicast & unknown unicast), but also known unicast frames. The egress PE determines whether or not to forward a particular frame to an Ethernet Segment depending on a combination of the split-horizon rule defined in [EVPN] and on the E-Tree

connectivity constraints:

- If the ESI Label indicates that the source Ethernet Segment is a Root, then the frame can be forwarded on a segment granted that it passes the split-horizon check.
- If the ESI Label indicates that the source Ethernet Segment is a Leaf, then the frame can be forwarded only on a Root segment, granted that it passes the split-horizon check.

When advertising the ESI MPLS label for a given Ethernet Segment, a PE must indicate whether the corresponding ESI is a Root or a Leaf site. This can be done by re-purposing one of the Reserved bits in the Flags field of the ESI MPLS label Extended Community attribute ([EVPN] Section 8) to indicate Root/Leaf status.

2.3 Scenario 3: Leaf AND Root site(s) per Ethernet Segment

In this scenario, a PE may have a set of one or more Root sites AND a set of one or more Leaf sites for a given VPN instance. An Ethernet Segment connected to the PE may be identified as both a Root and a Leaf site concurrently.

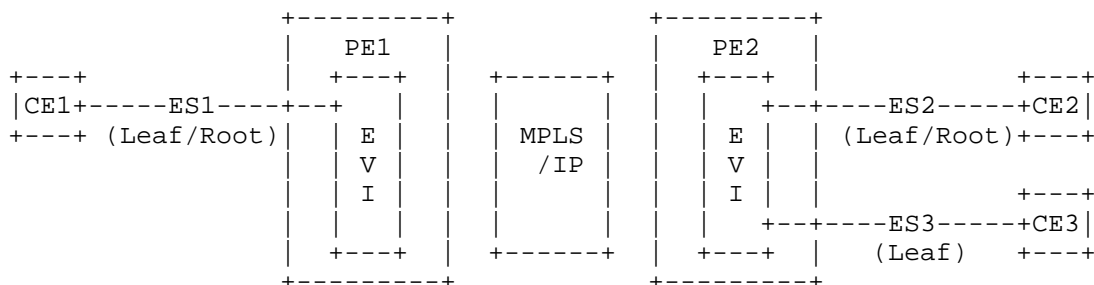


Figure 3: Scenario 3

This scenario can be addressed by extending the use of the ESI MPLS label, as described in the previous section, so that for an Ethernet Segment that has both Root and Leaf sites attached, two ESI MPLS labels are allocated and advertised: one ESI MPLS label denotes Root and the other denotes Leaf. The ingress PE imposes the right ESI MPLS label depending on whether the Ethernet frame originated from the Root or Leaf site on that Ethernet Segment. The mechanism by which the PE identifies whether a given frame originated from a Root or Leaf site on the segment is outside the scope of this document.

In addition to advertising two ESI MPLS labels per Ethernet Segment (for segments that have both Root and Leaf attached), a PE advertises

two special ESI MPLS labels: one for Root and another for Leaf. These are used by remote PEs for traffic originating from single-homed segments and for multi-homed segments that are not connected to the advertising PE.

3 Operation

Per [ETREE-FRAMEWORK], a generic E-Tree service supports all of the following traffic flows:

- Ethernet Unicast from Root to Leaf
- Ethernet Unicast from Leaf to Root
- Ethernet Unicast from Root to Root
- Ethernet Broadcast/Multicast from Root to Roots & Leafs
- Ethernet Broadcast/Multicast from Leaf to Roots

A particular E-Tree service may need to support all of the above types of flows or only a select subset, depending on the target application. In the case where unicast flows need not be supported, the L2VPN PEs can avoid performing any MAC learning function.

In the subsections that follow, we will describe the operation of E-VPN to support E-Tree service with and without MAC learning.

3.1 E-Tree with MAC Learning

The PEs implementing an E-Tree service must perform MAC learning when unicast traffic flows must be supported from Root to Leaf or from Leaf to Root sites. In this case, the PE with Root sites performs MAC learning in the data-path over the Ethernet Segments, and advertises reachability in E-VPN MAC Advertisement routes. These routes will be imported by PEs that have Leaf sites as well as by PEs that have Root sites, in a given EVI. Similarly, the PEs with Leaf sites perform MAC learning in the data-path over their Ethernet Segments, and advertise reachability in E-VPN MAC Advertisement routes which are imported only by PEs with at least one Root site in the EVI. A PE with only Leaf sites will not import these routes. PEs with Root and/or Leaf sites may use the Ethernet A-D routes for aliasing (in the case of multi-homed segments) and for mass MAC withdrawal.

To support multicast/broadcast from Root to Leaf sites, either a P2MP tree rooted at the PE(s) with the Root site(s) or ingress replication can be used. The multicast tunnels are set up through the exchange of the E-VPN Inclusive Multicast route, as defined in [E-VPN].

To support multicast/broadcast from Leaf to Root sites, ingress replication should be sufficient for most scenarios where there is a single Root or few Roots. If the number of Roots is large, a P2MP

tree rooted at the PEs with Leaf sites may be used.

3.2 E-Tree without MAC Learning

The PEs implementing an E-Tree service need not perform MAC learning when the traffic flows between Root and Leaf sites are multicast or broadcast. In this case, the PEs do not exchange E-VPN MAC Advertisement routes. Instead, the Ethernet A-D routes are used to exchange the E-VPN labels.

The fields of the Ethernet A-D route are populated per the procedures defined in [E-VPN], and the route import rules are as described in previous sections.

4 Acknowledgement

We would like to thank Sami Boutros for his comments.

5 Security Considerations

Same security considerations as [E-VPN].

6 IANA Considerations

Allocation of Extended Community Type and Sub-Type for E-VPN.

7 References

7.1 Normative References

[KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[ETREE-FRAMEWORK] Key et al., "A Framework for E-Tree Service over MPLS Network", draft-ietf-l2vpn-etree-frwk-00, work in progress, January 2012.

7.2 Informative References

[EVPN] Aggarwal et al., "BGP MPLS Based Ethernet VPN", draft-ietf-l2vpn-evpn-00.txt, work in progress, February, 2012.

[ETREE-REQ] Key et al., "Requirements for MEF E-Tree Support in VPLS", draft-ietf-l2vpn-etree-req-01.txt, work in progress, April, 2012.

Authors' Addresses

Ali Sajassi
Cisco
170 West Tasman Drive
San Jose, CA 95134, US
Email: sajassi@cisco.com

Samer Salam
Cisco
595 Burrard Street, Suite 2123
Vancouver, BC V7X 1J1, Canada
Email: ssalam@cisco.com