

Internet Working Group

Y. Jiang

Internet Draft

L. Yong
Huawei

Intended status: Standards Track

M. Paul
Deutsche Telekom

F. Jounay
Orange CH

F. Balus
W. Henderickx
Alcatel-Lucent

A. Sajassi
Cisco

Expires: August 2013

February 5, 2013

VPLS PE Model for E-Tree Support
draft-ietf-l2vpn-vpls-pe-etree-01.txt

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/ietf/lid-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 August 5, 2013.

Copyright Notice

Copyright (c) 2013 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.

Abstract

A generic VPLS solution for E-Tree services is proposed which uses VLANs to indicate root/leaf traffic. A VPLS Provider Edge (PE) model is illustrated as an example for the solution. In the solution, E-Tree VPLS PEs are interconnected by PWs which carry the VLAN indicating the E-Tree attribute, the MAC address based Ethernet forwarding engine and the PW work in the same way as before. A signaling mechanism for E-Tree capability and VLAN mapping negotiation is further described.

Table of Contents

1.	Introduction	3
2.	Conventions used in this document	4
3.	Terminology	4
4.	PE Model with E-Tree Support	5
4.1.	Existing PE Models	5
4.2.	A New PE Model with E-Tree Support	8
5.	PW for E-Tree Support	9
5.1.	PW Encapsulation	9
5.2.	VLAN Mapping	9
5.3.	PW Processing	11
5.3.1.	PW Processing in the VLAN Mapping Mode	11
5.3.2.	PW Processing in the Compatible Mode	12
5.3.3.	PW Processing in the Optimized Mode	13
6.	Signaling for E-Tree Support	14
6.1.	LDP Extensions for E-Tree Support	14
6.2.	BGP Extensions for E-Tree Support	16
7.	OAM Considerations	18
8.	Applicability	18
9.	Security Considerations	18

10. IANA Considerations	18
11. References	19
11.1. Normative References	19
11.2. Informative References	19
12. Acknowledgments	20
Appendix A. Other PE Models for E-Tree	21
A.1. A PE Model With a VSI and No bridge	21
A.2. A PE Model With external E-Tree interface	22

1. Introduction

The E-Tree service is defined in Metro Ethernet Forum (MEF) as a Rooted-Multipoint EVC service. It is a multipoint Ethernet service with special restrictions: the frames from a root may be received by any other root or leaf, and the frames from a leaf may be received by any root, but MUST not be received by a leaf. Further, an E-Tree service may include multiple roots and multiple leaves. Although VPMS or P2MP multicast is a somewhat simplified version of this service, in fact, there is no exact corresponding terminology in IETF.

[Etree-req] gives the requirements for providing E-Tree solutions in the VPLS and the need to filter leaf-to-leaf traffic.

[Vpls-etree] describes a PW control word based E-Tree solution, where a bit in the PW control word is used to indicate the root/leaf attribute for a packet. The Ethernet forwarder in the VPLS is also extended to filter the leaf-to-leaf traffic based on the <ingress port, egress port, CW L-bit> tuple.

[Etree-2PW] proposes another E-Tree solution where root and leaf traffic are classified and forwarded in the same VSI but with two separate PWs.

Both solutions are only applicable to "VPLS only" networks.

In fact, VPLS PE usually consists of a bridge module itself (see [RFC4664] and [RFC6246]); moreover, E-Tree services may cross both Ethernet and VPLS domains. Therefore, it is necessary to develop an E-Tree solution both for "VPLS only" scenarios and for interworking between Ethernet and VPLS.

IEEE 802.1 has incorporated the generic E-Tree solution in the latest version of 802.1Q [802.1aq], which is just an improvement on the traditional asymmetric VLAN mechanism (the use of different VLANs to indicate E-Tree root/leaf attributes and prohibiting leaf-to-leaf

traffic with the help of VLANs was first standardized in IEEE 802.1Q-2003). In the solution, VLANs are used to indicate root/leaf attribute of a packet: one VLAN ID is used to indicate the frames originated from the roots and another VLAN ID is used to indicate the frames originated from the leaves. At a leaf port, the bridge can then filter out all the frames from other leaf ports based on the VLAN ID. It is better to reuse the same mechanism in VPLS than to develop a new mechanism. The latter will introduce more complexity to interwork with IEEE 802.1Q solution.

This document introduces how the Ethernet VLAN solution can be used to support generic E-Tree services in the VPLS. The solution proposed here is fully compatible with the IEEE bridge architecture and the IETF PWE3 technology, thus it will not change the FIB (such as installing E-Tree attributes in the FIB), or need any specially tailored implementation. Furthermore, VPLS scalability and simplicity is also well kept. With this mechanism, it is also convenient to deploy a converged E-Tree service across both Ethernet and MPLS networks.

Firstly, a typical VPLS PE model is introduced as an example; the model is then extended in which a Tree VSI is connected to a VLAN bridge with a dual-VLAN interface.

This document then discusses the PW encapsulation and PW processing such as VLAN mapping options for transporting E-Tree services in a VPLS.

Finally, it describes the signaling extensions for E-Tree support and PE processing procedures.

2. 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 [RFC2119].

3. Terminology

E-Tree: a Rooted-Multipoint EVC service as defined in MEF 6.1

EVC: Ethernet Virtual Connection, as defined in MEF 4.0

FIB: Forwarding Information Base, or forwarding table

T-VSI: Tree VSI, a VSI with E-Tree support

Root AC, an AC attached with a root

Leaf AC, an AC attached with a leaf

C-VLAN, Customer VLAN

S-VLAN, Service VLAN

B-VLAN, Backbone VLAN

Root VLAN, a VLAN ID used to indicate all the frames that are originated at a root AC

Leaf VLAN, a VLAN ID used to indicate all the frames that are originated at a leaf AC

I-SID, Backbone Service Instance Identifier, as defined in IEEE 802.1ah

4. PE Model with E-Tree Support

"VPLS only" PE architecture as shown in Fig. 1 of [Etree-req] is a simplification of the VPLS and PWE3 architecture, several common VPLS PE architectures are discussed in more details in [RFC4664] and [RFC6246].

Therefore, VLAN based E-Tree solution are demonstrated with the help of a typical VPLS PE model. It can also be used by other PE models which are discussed in Appendix A.

4.1. Existing PE Models

According to [RFC4664], there are at least three models possible for a VPLS PE, including:

- o A single bridge module, a single VSI;
- o A single bridge module, multiple VSIs;
- o Multiple bridge modules, each attaches to a VSI.

The second PE model is commonly used. A typical example is further depicted in Fig. 1 and Fig. 2 [RFC6246], where an S-VLAN bridge

module is connected to multiple VSIs each with a single VLAN virtual interface.




Figure 1 A model of 802.1ad Bridge Module




Figure 2 A VPLS-capable PE Model

In this PE model, Ethernet frames from Customer Edges (CEs) will cross multiple stages of bridge modules (i.e., C-VLAN and S-VLAN

bridge) and a VSI in a PE before being sent on the PW to a remote PE. Therefore, the association between an AC port and a PW on a VSI as required in [Vpls-etree] or [Etree-2PW] is difficult, sometimes even impossible.

This model could be further enhanced: When Ethernet frames arrive at a PE, a root VLAN or a leaf VLAN tag is added. Then the frames with the root VLAN tag are transmitted both to the roots and the leaves, while the frames with the leaf VLAN tag are transmitted to the roots but dropped for the leaves (these VLAN tags are removed before the frames are transmitted over the wire). It was demonstrated in [802.1aq] that the E-Tree service in Ethernet networks can be well supported with this mechanism.

Assuming this mechanism is implemented in the bridge module, it is quite straightforward to infer a VPLS PE model with two VSIs to support the E-Tree (as shown in Fig. 3). But this model will require two VSIs per PE and two sets of PWs per E-Tree service, which is poorly scalable in a large MPLS/VPLS network; in addition, both these VSIs have to share their learned MAC addresses.

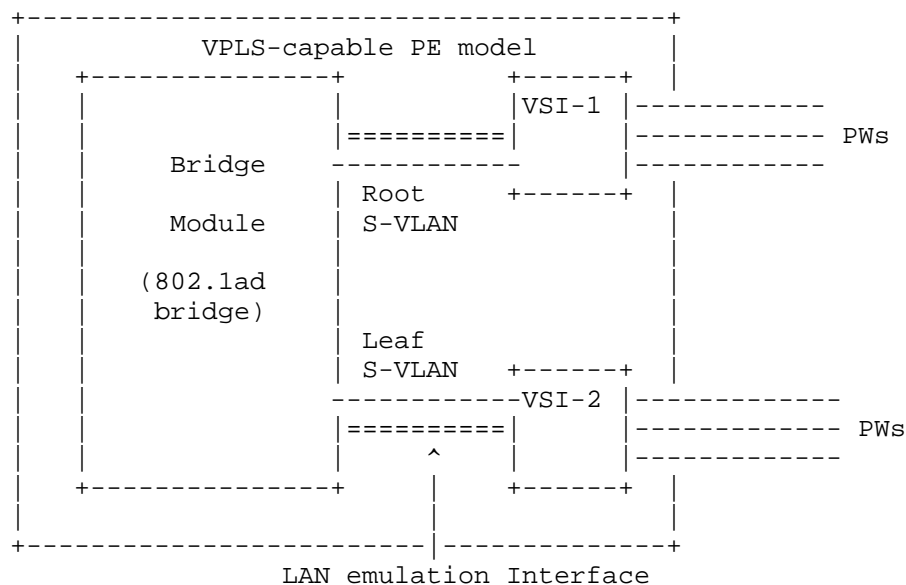


Figure 3 A VPLS PE Model for E-Tree with 2 VSIs

4.2. A New PE Model with E-Tree Support

In order to support the E-Tree in a more scalable way, a new VPLS PE model with a single Tree VSI (T-VSI, a VSI with E-Tree support) is proposed. As depicted in Fig. 4, the bridge module is connected to the T-VSI with a dual-VLAN virtual interface, i.e., both the root VLAN and the leaf VLAN are connected to the same T-VSI, and they share the same FIB and work in shared VLAN learning. In this way, only one VPLS instance and one set of PWs is needed per E-Tree service, and the scalability of VPLS is improved.

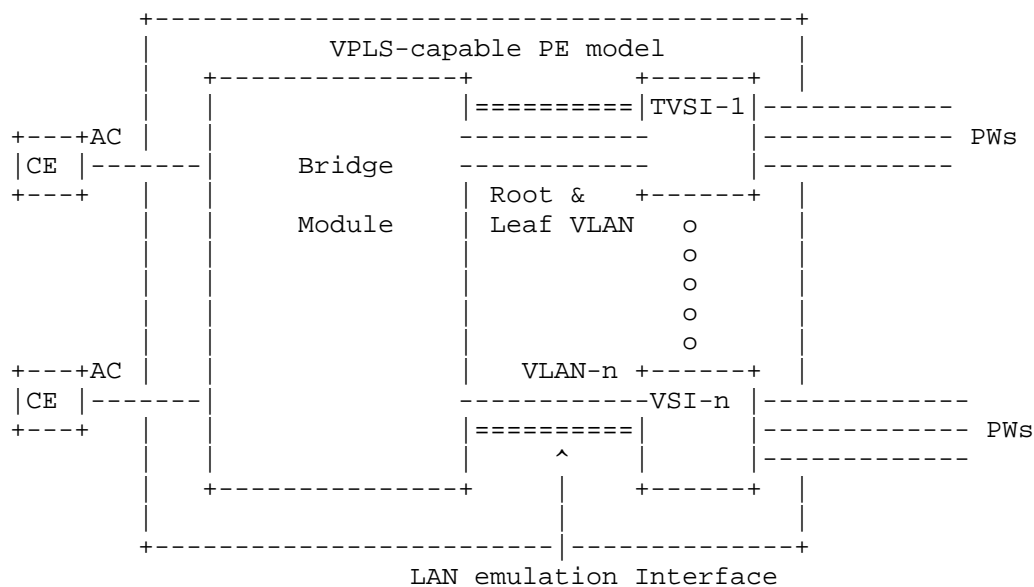


Figure 4 A VPLS PE Model for E-Tree with a Single T-VSI

For an untagged port (frames over this port are untagged) or VLAN-unaware port (VLAN tags in the frames are ignored), the Ethernet frames received from the root ACs SHOULD be tagged with a root C-VLAN, and optionally MAY be added with another root S-VLAN.

For a C-VLAN tagged port, the Ethernet frames received from the root ACs SHOULD be added with a root S-VLAN.

For an S-VLAN tagged port, the S-VLAN tag in the Ethernet frames received from the root ACs SHOULD be translated to the root S-VLAN in the VPLS network domain. Alternatively, the PBB VPLS PE model (where an IEEE 802.1ah bridge module is embedded in the PE) as described in [PBB-VPLS] MAY be used, and a root B-VLAN or leaf B-VLAN MAY be added

in this case (the E-Tree attribute may also be indicated with two I-SID tags in the bridge module, and the frames are further encapsulated and transported transparently over a single B-VLAN, thus the PBB VPLS works just in the same way as described in [PBB-VPLS] and will be discussed no more in this document). When many S-VLANs are multiplexed in a single AC, the 2nd option has an advantage of both VLAN scalability and MAC address scalability.

In a similar way, the traffic from the leaf ACs is tagged and transported on the leaf C-VLAN, S-VLAN or B-VLAN.

In all cases, the outermost VLAN in the resulted Ethernet header is used to indicate the E-Tree attribute of an Ethernet frame; this document will use VLAN to refer to this outermost VLAN for simplicity in the latter sections.

5. PW for E-Tree Support

5.1. PW Encapsulation

To support an E-Tree service, T-VSIs in a VPLS must be interconnected with a bidirectional Ethernet PW. The Ethernet PW may work in the tagged mode (PW type 0x0004) as described in [RFC4448], and a VLAN tag must be carried in each frame in the PW to indicate the frame originated from either root or leaf (the VLAN tag indicating the frame originated from either root or leaf can be translated by a bridge module in the PE or added by an outside Ethernet edge device, even by a customer device). In the tagged PW mode, two service delimiting VLANs must be allocated in the VPLS domain for an E-Tree. PW processing for the tagged PW will be described in Section 5.3 of this document.

Raw PW (PW type 0x0005 in [RFC4448]) may be used to carry E-Tree service for a PW in Compatible mode as shown in Section 5.3.2.

5.2. VLAN Mapping

There are two ways of manipulating VLANs for an E-Tree in VPLS:

- o Global VLAN based, that is, provisioning two global VLANs (Root VLAN, Leaf VLAN) across the VPLS network, thus no VLAN mapping is needed at all, or the VLAN mapping is done completely in the Ethernet domains.

- o Local VLAN based, that is, provisioning two local VLANs for each PE (which participates in the E-Tree) in the VPLS network independently.

The first method requires no VLAN mapping in the PW, but two unique service delimiting VLANs must be allocated across the VPLS domain.

The second method is more scalable in the use of VLANs, but needs a VLAN mapping mechanism in the PW similar to what is already described in Section 4.3 of [RFC4448].

Global or local VLANs can be manually configured or provisioned by an OSS system. Alternatively, some automatic VLAN allocation algorithm may be provided in the management plane, but it is out scope of this document.

For both methods, VLAN mapping parameters from a remote PE can be provisioned or determined by a signaling protocol as described in Section 6 when a PW is being established.

5.3. PW Processing

5.3.1. PW Processing in the VLAN Mapping Mode

In the VLAN Mapping mode, two VPLS PE with E-Tree capability are inter-connected with a PW (For example, the scenario of Fig. 5 depicts the interconnection of two PEs miscellaneously attached with both root and leaf nodes).

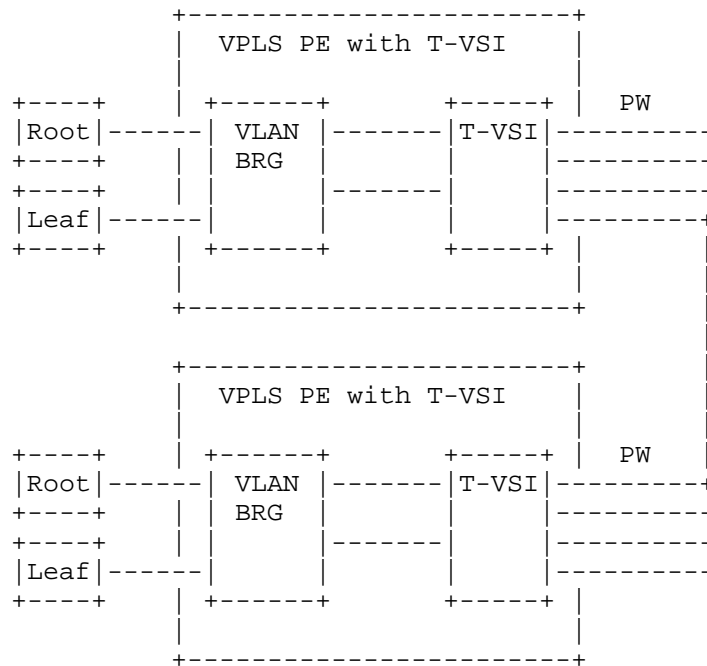


Figure 5 T-VSI Interconnected in the Normal Mode

If a PE is in the VLAN mapping mode for a PW, then in the data plane the PE MUST map the VLAN in each frame as follows:

- o Upon transmitting frames on the PW, map from local VLAN to remote VLAN (i.e., the local leaf VLAN in a frame is translated to the remote leaf VLAN; the local root VLAN in a frame is translated to the remote root VLAN).
- o Upon receiving frames on the PW, map from remote VLAN to local VLAN, and the frames are further forwarded or dropped in the egress bridge module using the filtering mechanism as described in [802.1aq].

The signaling for VLANs is specified in Section 6.

5.3.2. PW Processing in the Compatible Mode

The new VPLS PE model can work in a traditional VPLS network seamlessly in the compatibility mode. As shown in Fig. 6, the VPLS PE with T-VSI can be attached with root and/or leaf nodes, while the VPLS PE with a traditional VSI can only be attached with root nodes. A raw PW should be used to connect them.




Figure 6 T-VSI interconnected with Traditional VSI

If a PE is in the Compatible mode for a PW, then in the data plane the PE MUST process the frame as follows:

- o Upon transmitting frames on the PW, remove the root or leaf VLAN in the frames.
- o Upon receiving frames on the PW, add a VLAN tag with a value of the local root VLAN to the frames.

5.3.3.PW Processing in the Optimized Mode

When two PEs (both have E-Tree capability) are inter-connected and one of them (e.g., PE2) is attached with only leaf nodes, as shown in the scenario of Fig. 7, its peer PE (e.g., PE1) should then work in the optimized mode. In this case, PE1 should not send the frames originated from the local leaf VLAN to PE2, i.e., these frames are dropped rather than transported over the PW. The bandwidth efficiency of the VPLS can thus be improved. The signaling for the PE attached with only leaf nodes is specified in Section 6.




Figure 7 T-VSI interconnected with PE attached with only leaf nodes

If a PE is in the Optimized Mode for a PW, upon transmit, the PE SHOULD first operate as follows:

- o Drop a frame if its VLAN ID matches the local leaf VLAN ID.

6. Signaling for E-Tree Support

6.1. LDP Extensions for E-Tree Support

In addition to the signaling procedures as specified in [RFC4447], this document proposes a new interface parameter sub-TLV to provision an E-Tree service and negotiate the VLAN mapping function, as follows:




Figure 8 E-Tree Sub-TLV

Where:

- o E-Tree is the sub-TLV identifier to be assigned by IANA.
- o Length is the length of the sub TLV in octets.
- o Reserved bits MUST be set to zero on transmit and be ignored on receive.
- o P is a Leaf-only bit, it is set to 1 to indicate that the PE is attached with only leaf nodes, and set to 0 otherwise.
- o V is a bit indicating the sender's VLAN mapping capability. A PE capable of VLAN mapping MUST set this bit, and clear it otherwise.
- o Root VLAN ID is the value of the local root VLAN.
- o Leaf VLAN ID is the value of the local leaf VLAN.

When setting up a PW for the E-Tree based VPLS, two PEs negotiate the E-Tree support using the above E-Tree sub-TLV. Note PW type of 0x0004 should be used during the PW negotiation.

A PE that wishes to support E-Tree service MUST include an E-Tree Sub-TLV in its PW label mapping message and include its local root VLAN ID and leaf VLAN ID in the TLV. A PE that has the VLAN mapping capability MUST set the V bit to 1, and a PE is attached with only leaf nodes SHOULD set the P bit to 1.

In default, for each PW, VLAN-Mapping-Mode, Compatible-Mode, and Optimized-Mode are all set to FALSE.

A PE that receives a PW label mapping message with an E-Tree Sub-TLV from its peer PE, after saving the VLAN information for the PW, must process it as follows:

- 1) if the root and leaf VLAN ID in the message match the local root and leaf VLAN ID, then continue to 3);
 - 2) else {
 - if the bit V is cleared, then {
 - if the PE is capable of VLAN mapping, then it MUST set VLAN-Mapping-Mode to TRUE;
 - else {
 - A label release message with the error code "E-Tree VLAN mapping not supported" is sent to the peer PE and exit the process;
 - if the bit V is set, and the PE is capable of VLAN mapping, then the PE with the minimum IP address MUST set VLAN-Mapping-Mode to TRUE;
- 3) If the P bit is set, then:
 - {
 - If the PE is a leaf-only node itself, then a label release message with a status code "Leaf to Leaf PW released" is sent to the peer PE and exit the process;
 - Else the PE SHOULD set the Optimized-Mode to TRUE.

If a PE has sent an E-Tree Sub-TLV but does not receive any E-Tree Sub-TLV in its peer's PW label mapping message, The PE SHOULD then

establish a raw PW with this peer as in traditional VPLS and set Compatible-Mode to TRUE for this PW.

Data plane processing for this PW is as following:

If Optimized-Mode is TRUE, then data plane processing as described in Section 5.3.3 applies.

If VLAN-Mapping-Mode is TRUE, then data plane processing as described in Section 5.3.1 applies.

If Compatible-Mode is TRUE, then data plane processing is as described in Section 5.3.2.

PW processing as described in [RFC4448] proceeds as usual for all cases.

6.2. BGP Extensions for E-Tree Support

A new E-Tree extended community is proposed for E-Tree signaling in BGP VPLS:




Figure 9 E-Tree Extended Community

Where:

- o Root VLAN ID is the value of the local root VLAN.
- o Leaf VLAN ID is the value of the local leaf VLAN.
- o Reserved, 15 bits MUST be set to zero on transmit and be ignored on receive.
- o P is a Leaf-only bit, it is set to 1 to indicate that the PE is attached with only leaf nodes, and set to 0 otherwise.

The PEs attached with both leaf and root nodes must support BGP E-Tree signaling as described in this document, and must support VLAN mapping in their data planes. The traditional PE attached with only root nodes may also participate in an E-Tree service.

In BGP VPLS signaling, besides attaching a Layer2 Info Extended Community as detailed in [RFC4761], an E-Tree Extended Community MUST be further attached if a PE wishes to participate in an E-Tree service. The PE MUST include its local root VLAN ID and leaf VLAN ID in the E-Tree Extended Community. A PE attached with only leaf nodes of an E-Tree SHOULD set the P bit in the E-Tree Extended Community to 1.

A PE that receives a BGP UPDATE message with an E-Tree Extended Community from its peer PE, after saving the VLAN information for the PW, must process it as follows (after processing procedures as specified in Section 3.2 of [RFC4761]):

- 1) if the root and leaf VLAN ID in the E-Tree Extended Community match the local root and leaf VLAN ID, then continue to 3);
- 2) else {
 the PE with the minimum IP address MUST set VLAN-Mapping-Mode to TRUE;
}
- 3) If the P bit is set {
 If the PE is a leaf-only PE itself, then forbids any traffic on the PW;

 Else the PE SHOULD set the Optimized-Mode to TRUE.
}

A PE which does not recognize this attribute shall ignore it silently. If a PE has sent an E-Tree Extended Community but does not receive any E-Tree Extended Community from its peer, the PE SHOULD then establish a raw PW with this peer as in traditional VPLS, and set Compatible-Mode to TRUE for this PW.

Data plane in the VPLS is the same as described in Section 4.2 of [RFC4761], and data plane processing for a PW is the same as described at the end of Section 6.1.

7. OAM Considerations

VPLS OAM requirements and framework as specified in [RFC6136] are applicable to E-Tree, as both Ethernet OAM frames and data traffic are transported over the same PW.

Ethernet OAM for E-Tree including both service OAM and segment OAM frames shall undergo the same VLAN mapping as the data traffic; and root VLAN SHOULD be applied to segment OAM frames so that they are not filtered.

8. Applicability

The solution is applicable to both LDP VPLS [RFC4762] and BGP VPLS [RFC4761].

The solution is applicable to both "VPLS Only" networks and VPLS with Ethernet aggregation networks.

The solution is also applicable to PBB VPLS networks.

9. Security Considerations

Besides security considerations as described in [RFC4448], [RFC4761] and [RFC4762], this solution prevents leaf to leaf communication in the data plane of VPLS when its PEs are interconnected with PWs. In this regard, security can be enhanced for customers with this solution.

10. IANA Considerations

IANA is requested to allocate a value for E-Tree in the registry of Pseudowire Interface Parameters Sub-TLV type.

Parameter ID	Length	Description
TBD	8	E-Tree

IANA is requested to allocate two new LDP status codes from the registry of name "STATUS CODE NAME SPACE". The following values are suggested:

Range/Value	E	Description
-----	----	-----
TBD	1	E-Tree VLAN mapping not supported
TBD	0	Leaf to Leaf PW released

IANA is requested to allocate a value for E-Tree in the registry of BGP Extended Community.

Type Value	Name
=====	=====
TBD	E-Tree Info

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4447] Martini, L., and et al, "Pseudowire Setup and Maintenance Using Label Distribution Protocol (LDP)", RFC 4447, April 2006.
- [RFC4448] Martini, L., and et al, "Encapsulation Methods for Transport of Ethernet over MPLS Networks", RFC 4448, April 2006.
- [RFC4761] Kompella, K. and Rekhter, Y., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007
- [RFC4762] Lasserre, M. and Kompella, V., "Virtual Private LAN Services using LDP", RFC 4762, January 2007.
- [RFC6136] Sajassi, A. and Mohan, D., "L2VPN OAM Requirements and Framework", RFC 6136, March 2011

11.2. Informative References

- [RFC3985] Bryant, S., and Pate, P., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.
- [RFC4664] Andersson, L., and Rosen, E., "Framework for Layer 2 Virtual Private Networks (L2VPNs)", RFC 4664, September 2006.

- [RFC6246] Sajassi, A., and et al, "Virtual Private LAN Service (VPLS) Interoperability with Customer Edge (CE) Bridges", RFC 6246, June 2011
- [ETree-req] Key, R., et al, "Requirements for MEF E-Tree Support in VPLS", draft-ietf-l2vpn-etree-reqt-01, Work in Progress
- [Vpls-etree] Key, R., and et al, "Extension to VPLS for E-Tree", draft-key-l2vpn-vpls-etree-06, October 2011
- [802.1aq] IEEE 802.1aq D4.3, Virtual Bridged Local Area Networks - Amendment 9: Shortest Path Bridging, September 2011
- [Etree-2PW] Ram, R., and et al., Extension to LDP-VPLS for E-Tree Using Two PW, draft-ram-l2vpn-ldp-vpls-etree-2pw-02, May 2011
- [PBB-VPLS] Balus, F., and et al., Extensions to VPLS PE model for Provider Backbone Bridging, draft-ietf-l2vpn-pbb-vpls-pe-model-04, October 2011

12. Acknowledgments

The authors would like to thank Adrian Farrel, Susan Hares and Shane Amante for their valuable advices, thank Ben Mack-crane, Edwin Mallette, Donald Fedyk, Dave Allan, Giles Heron, Raymond Key, Josh Rogers, Sam Cao and Daniel Cohn for their valuable comments and discussions.

Appendix A. Other PE Models for E-Tree

A.1. A PE Model With a VSI and No bridge

If there is no bridge module in a PE, the PE may consist of Native Service Processors (NSPs) as shown in Figure A.1 (adapted from Fig. 5 of [RFC3985]) where any transformation operation for VLANs (e.g., VLAN insertion/removal or VLAN mapping) may be applied. Thus a root VLAN or leaf VLAN can be added by the NSP depending on the UNI type (root/leaf) associated with the AC over which the packet arrives.

Further, when a packet with a leaf VLAN exits a forwarder and arrives at the NSP, the NSP must drop the packet if the egress AC is associated with a leaf UNI.

Tagged PW and VLAN mapping work in the same way as in the typical PE model.




Figure A.1 A PE model with a VSI and no bridge module

This PE model may be used by an MTU-s in an H-VPLS network, or an N-PE in an H-VPLS network with non-bridging edge devices, wherein a spoke PW can be treated as an AC in this model.

A.2. A PE Model With external E-Tree interface




Figure A.2 A PE model with external E-Tree interface

A more simplified PE model is depicted in A.2, where Root/Leaf VLANs are directly or indirectly over a single PW connected to a same VSI forwarder in a PE, any transformation of E-Tree VLANs, e.g., VLAN insertion/removal or VLAN mapping, can be performed by some outer equipments, and the PE may further translate these VLANs into its own local VLANs. This PE model may be used by an N-PE in an H-VPLS network with bridging-capable devices, or scenarios such as providing E-Tree Network-to-Network (NNI) interfaces.

Authors' Addresses

Yuanlong Jiang
Huawei Technologies Co., Ltd.
Bantian, Longgang district
Shenzhen 518129, China
Email: jiangyuanlong@huawei.com

Lucy Yong
Huawei USA
1700 Alma Dr. Suite 500
Plano, TX 75075, USA
Email: lucyyong@huawei.com

Manuel Paul
Deutsche Telekom
Winterfeldtstr. 21
10781 Berlin, Germany
Email: manuel.paul@telekom.de

Frederic Jounay
Orange CH
4 rue caudray 1020 Renens, Switzerland
Email: frederic.jounay@orange.ch

Florin Balus
Alcatel-Lucent
701 E. Middlefield Road
Mountain View, CA, USA 94043
Email: florin.balus@alcatel-lucent.com

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

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

