

Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 15, 2011

R. Chen
Y. Wang
ZTE Corporation
L. Guo
China Telecom
March 14, 2011

**LDP Extensions for Optimized MAC Address Withdrawal in VPLS model 3
supporting Qualified learning
draft-chen-l2vpn-vpls-mac-opt-qualified-01**

Abstract

This draft defines a method for optimized MAC Address Withdrawal in VPLS model 3 ([[RFC4664](#)]) supporting qualified learning ([[RFC4762](#)]). Some extensions are made based on the MAC Address Withdrawal procedures defined in [[RFC4762](#)] and [[I-D.ietf-l2vpn-vpls-ldp-mac-opt](#)], in order to enable a PE to remove only the MAC addresses, which belong to MAC address space affected by topology change.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on September 15, 2011.

Copyright Notice

Copyright (c) 2011 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
2.	Conventions used in this document	3
2.1.	Terminology	4
3.	Problem Description	5
4.	Optimized MAC Flush Mechanism	6
4.1.	MAC Address Space TLV	6
4.2.	MAC Address Space TLV Processing Rules	7
4.3.	Optimized MAC Flush Procedures	8
5.	IANA Considerations	8
6.	Security Considerations	8
7.	Normative references	8
	Authors' Addresses	9

1. Introduction

There are operators deploying VPLS with qualified learning mode [[RFC4762](#)], so as to reduce the number of VPLS instance. By using VLAN within VPLS and qualified learning mode to separate customer internally, current defined VPLS MAC address withdraw mechanism will flush MAC addresses VLAN space which are not affected due to topology change.

[RFC4762] defines a basic MAC Address Withdrawal mechanism to remove or unlearn MAC addresses for faster convergence on topology change. It defines MAC List TLV which contains a list of MAC addresses to be flushed in LDP Address Withdraw Message, and when a MAC List TLV contains a large number of MAC addresses, it would be preferable to send a LDP Address Withdraw Message with an empty MAC List.

As per the rules in [[RFC4762](#)], on receiving a MAC Address Withdrawal Message with MAC List TLV, a PE removes the association between the MAC address and the AC or PW over which this message is received. For a MAC Address Withdraw message with empty MAC list, a PE removes all MAC addresses except the MAC addresses learned over the newly activated PW. The PE further transmits a MAC Address Withdrawal message to each remote PE in the VPLS network. The problem is that it will flush MAC addresses which are not affected due to topology change, thus leading to unnecessary flooding and relearning.

Draft [[I-D.ietf-l2vpn-vpls-ldp-mac-opt](#)] defines an extension to the MAC Address Withdrawal procedure with empty MAC List ([[RFC4762](#)]). It is accomplished by sending an LDP Address Withdraw Message with the empty MAC addresses and PE-ID TLV. When a PE device receives a MAC Address Withdrawal message with a PE-ID TLV, it SHOULD flush all the MAC addresses learned from the PW, which terminates in the remote VSI identified by the PE-ID element. Comparing with [[RFC4762](#)], this mechanism narrows the scope of MAC address flush to PE-ID elements. But if this mechanism is used in VPLS model 3 ([[RFC4664](#)]) supporting qualified learning ([[RFC4762](#)]), it also flushes MAC addresses belonging to MAC address space which are not affected due to topology change.

This draft describes the problem and a solution to optimize the MAC flush procedure in [[RFC4762](#)] and [[I-D.ietf-l2vpn-vpls-ldp-mac-opt](#)], and enable it to flush only the MAC addresses belonging to the MAC address space affected by topology change.

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](#).

2.1. Terminology

AC: Attachment Circuit

CE: Customer Edge device

LDP: Label Distribution Protocol

PE: Provider Edge

PW: Pseudowire

MTU-s: Multi-Tenant Unit switch

3. Problem Description

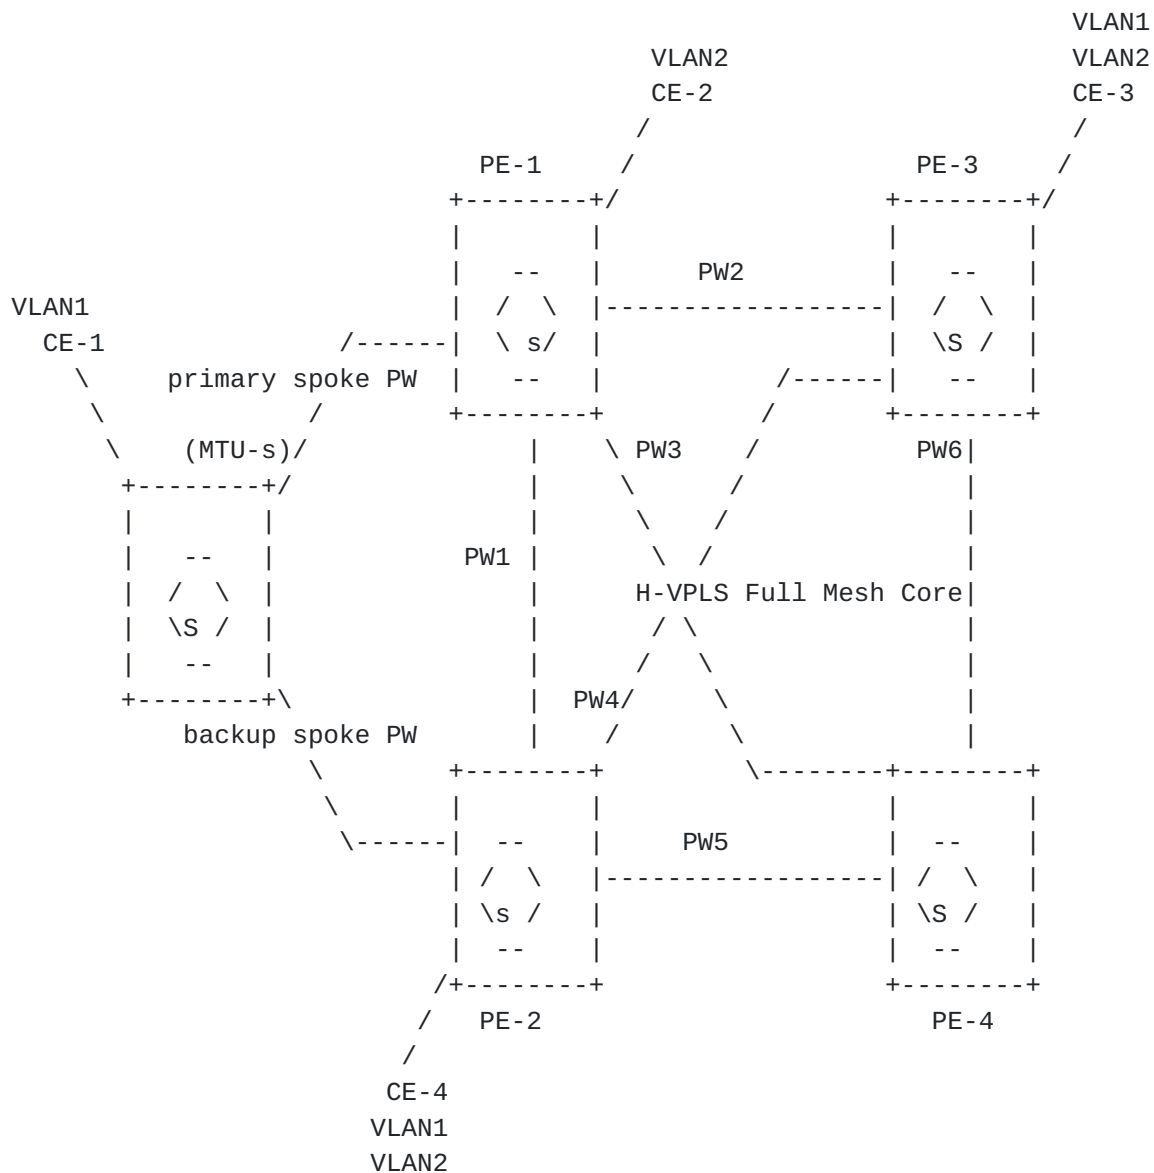


Figure 1: Dual homed MTU-s in two tier hierarchy H-VPLS

Figure 1 is an example for the dual-homing access topology in VPLS model 3 supporting qualified learning.

There are four VLANs in Figure 1, and each customer VLAN has its own broadcast domain and MAC address space. CE-1 and CE-3 belong to VLAN1, CE-2 and CE-4 belong to VLAN2.

CE-1, CE-2, CE-3 and CE-4 connect with MTU-S, PE-1, PE-4 and PE-2

respectively. The MTU-S is dual-homed to PE-1 and PE-2. Only the primary spoke PW is active at MTU-s, thus PE-1 is acting as the active device to reach the full mesh in the VPLS instance.

When MTU-s switches to the backup spoke PW and activates it, PE-2 becomes the active device to reach the full mesh core. Traffic entering the H-VPLS from CE-1 is diverted by the MTU-s to the backup spoke PW. For faster convergence MTU-s may desire to unlearn or remove the MAC addresses belonging to VLAN1 and learned from the PW that terminates at PE-1 MUST be removed. Once the backup PW has been made active, MTU-s may send a MAC flush message to PE-2.

As per the rules defined in [[RFC4762](#)], PE-2 flushes all of the MAC addresses learned in the VPLS from the PWs terminating at PE-1, PE-3 and PE-4. PE-2 further relays MAC flush messages to PE-1, PE-3 and PE-4. Same processing rule applies at all those PE.

As per the rules defined in [[I-D.ietf-l2vpn-vpls-ldp-mac-opt](#)], PE-2 flushes all of the MAC addresses learned in the VPLS from the PWs terminating at PE-1. In fact, PE2 need not flush the MAC addresses belonging to VLAN2 and learned over PW1. There are multi-Mac addresses spaces in a single VPLS, and the affected Mac addresses spaces is only VLAN1. PE-2 only need flush and relearn MAC addresses belong to VLAN1 and learned over PW1.

With the number of PE in the full-mesh increases, the number of unaffected MAC addresses flushed in a VPLS instance also increases, thus leading to unnecessary flooding and relearning.

4. Optimized MAC Flush Mechanism

4.1. MAC Address Space TLV

This draft proposes a MAC Address Space TLV, it is carried in the LDP Address Withdraw Message. The MAC Address Space TLV carries a MAC address Space which are affected due to topology change. When a PE receives a LDP Address Withdraw message with MAC Address Space TLV, then the PE only flushes the suspect MAC addresses belonging to the MAC Address Space elements.

The new MAC Address Space TLV format is as follows:




Figure 2: MAC Address Space TLV format

U bit: Unknown bit. This bit MUST be set to 1. If the MAC address Space format is not understood, then it MUST be ignored.

F bit: Forward bit. This bit MUST be set to 0. Since the LDP mechanism used here is targeted, the TLV MUST NOT be forwarded.

Type: Type field. This identifies the TLV type as MAC Address Space TLV.

Length: Length field. This field specifies the total length in octets of value in MAC Address Space TLV.

MAC Address Space: This identifies the MAC addresses Space which are affected due to topology change.

The MAC address space TLV should immediately be after the MAC address list TLV. It is used to indicate MAC address space that these MAC address list belongs to.

4.2. MAC Address Space TLV Processing Rules

If a PE receives a MAC Address Withdraw with MAC Address Space TLV and empty MAC list, it SHOULD remove all the MAC addresses belonging to the MAC Address Spaces in a VPLS instance, except the MAC addresses learned over the PW associated with this signaling session over which the message was received.

If a PE receives a MAC Address Withdraw with MAC Address Space TLV and MAC list TLV, it SHOULD remove the MAC address belonging to the MAC Address Space in a VPLS instance.

If a PE receives a MAC Address Withdraw with MAC Address Space TLV and PE-ID TLV, from the specified MAC Address Space which is identified by MAC Address Space TLV, it SHOULD remove all the MAC addresses learned from the PW, which terminates at the PE identified by the PE-ID.

If a PE that doesn't support MAC Address Space TLV, receives a MAC flush message with this option, it MUST ignore the option and follow the processing rules as per [[RFC4762](#)] and [[I-D.ietf-l2vpn-vpls-ldp-mac-opt](#)].

4.3. Optimized MAC Flush Procedures

This section explains the optimized MAC flush procedure in the scenario shown in Figure 1.

When the backup PW is activated by MTU-s, it may send MAC Address Withdraw message to PE-2 with the optional PE-ID TLV and MAC Address Space TLV. Upon receipt of the MAC Address Withdraw message, PE-2 identifies the VPLS instance that requires MAC Address Withdraw from the FEC element in the FEC TLV. From the PE-ID TLV, PE-2 identifies the PW in the VPLS that terminates in PE-1. From the MAC Address Space TLV, PE-2 identifies MAC addresses Space which are affected due to topology change. PE-2 removes all MAC addresses belong to VLAN1 and learned over PW1.

PE-2 relays MAC Address Withdraw messages with the received MAC Address Space TLV and PE-ID to all its peer PE devices. When the message is received at PE-3/PE-4, MAC addresses belonging to VLAN1 and learned over the PW that terminates at PE1 MUST be removed.

5. IANA Considerations

The type field in the MAC List TLV is defined as 0x406 and is subject to IANA approval

6. Security Considerations

Security considerations discussed in [[RFC4762](#)] and [[I-D.ietf-l2vpn-vpls-ldp-mac-opt](#)] apply to this document.

7. Normative references

[[I-D.ietf-l2vpn-vpls-ldp-mac-opt](#)]
Dutta, P., Balus, F., Calvignac, G., and O. Stokes, "LDP Extensions for Optimized MAC Address Withdrawal in H-VPLS", [draft-ietf-l2vpn-vpls-ldp-mac-opt-03](#) (work in progress), October 2010.

[[RFC2119](#)] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [RFC4447] Martini, L., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", [RFC 4447](#), April 2006.
- [RFC4664] Andersson, L. and E. Rosen, "Framework for Layer 2 Virtual Private Networks (L2VPNs)", [RFC 4664](#), September 2006.
- [RFC4762] Lasserre, M. and V. Kompella, "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC 4762](#), January 2007.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.

Authors' Addresses

Ran Chen
ZTE Corporation
4F,RD Building 2,Zijinghua Road
Yuhuatai District,Nanjing 210012
P.R.China

Phone: +86 025 52878135
Email: chen.ran@zte.com.cn

Yubao Wang
ZTE Corporation
4F,RD Building 2,Zijinghua Road
Yuhuatai District,Nanjing 210012
P.R.China

Phone: +86 025 52872130
Email: wang.yubao2@zte.com.cn

Liang Guo
China Telecom
26/F,109 Zhongshan Ave
Tianhe District,Guangzhou 510630
P.R.China

Phone: +86 020 38639595
Email: guoliang@gsta.com

