

INTERNET-DRAFT
Intended Status: Informational

Mingui Zhang
Bin Wang
Liang Xia
Huawei
Jie Hu
China Telecom
August 12, 2014

Expires: February 13, 2015

Tagging Customer Bridge Domains in VPLS
draft-zhang-l2vpn-vpls-bd-tagging-02.txt

Abstract

This document proposes to use Customer VLAN ID as an identifier for traffic isolation in Virtual Private LAN Service (VPLS). In this way, multiple bridge domains of customers can share a single VPLS instance while their traffic are separated. With this proposal, Service Providers can be relieved from the heavy provisioning overhead of large number of pseudowires in the environment where a mass of bridge domains need be connected.

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) 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
2.	Acronyms and Terminology	3
2.1.	Acronyms	3
2.2.	Terminology	3
3.	PE Model	3
4.	Use Cases of U-tag Awareness in VPLS	4
4.1.	No Duplicated MAC Address	4
4.2.	Scalable Interconnection of L2 Sites	5
4.3.	BUM Traffic Scoped per BD	5
4.3.1.	Advertising Interested VLANs in LDP	5
4.3.2.	Dynamic VLAN Registration with MVRP	5
4.4.	Per C-VLAN MAC Withdraw	6
5.	Backward Compatibility	6
6.	Contributors	6
7.	Security Considerations	6
8.	IANA Considerations	6
9.	References	7
9.1.	Normative References	7
9.2.	Informative References	7
	Author's Addresses	8

1. Introduction

VPLS has been widely used to connect customers' bridge domains. Traffic segregation for customers is performed on a per VPLS instance basis. In the environment (e.g., Data Center Network) where a mass of customers multiplied with a plenty of bridge domains are to be connected, a large number of PWs need be maintained. Service Providers are therefore suffering from scalability issue.

This proposal suggests the Customer VLAN ID (U-tag) is used as an additional de-multiplexor for traffic segregation in VPLS. By doing this, multiple BDs can share the same VPLS instance while their traffic are isolated. This method can greatly reduce the number of PWs therefore reduce the provisioning overhead for operators. Use cases of this method are given in the document.

Two options arising from the industry are covered in the discussion. The first one is proposed in [[V-aware](#)]. It extends the LDP control plane for PEs to advertise supported VLANs. The second option makes use of VLAN registration protocol, such as [[MVRP](#)], to exchange supported C-VLANs between PEs.

2. Acronyms and Terminology

2.1. Acronyms

MVRP: Multiple VLAN Registration Protocol
BD: Bridge Domain/Broadcast Domain
PW: Pseudowire
VSI: Virtual Switch Instance
U-tag: Customer VLAN ID
C-VLAN: Customer VLAN
BUM: Broadcast, Unknown unicast and Multicast
VLL: Virtual Leased Line

2.2. 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](#)].

3. PE Model

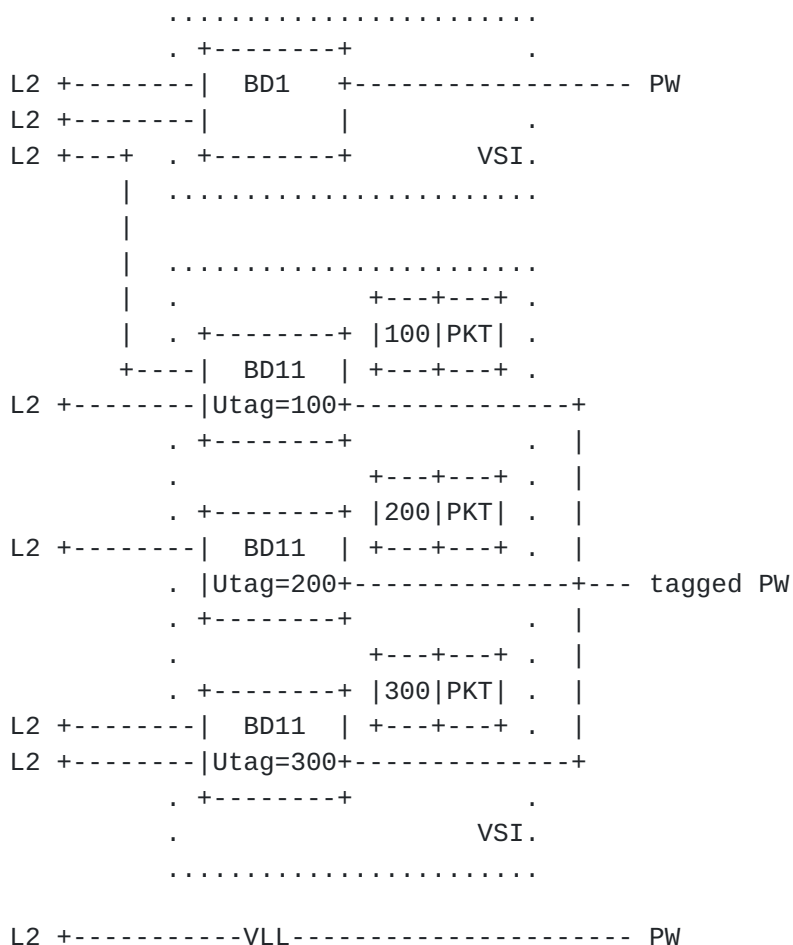


Figure 3.1: U-tag is used as the service de-multiplexor in tagged PW

In Figure 3.1, an example is used to show that the Customer VLAN ID (U-tag) is used as a finer grained de-multiplexor for traffic segregation. Therefore, multiple customer BDs can be integrated into one VSI while their traffic is isolated.

4. Use Cases of U-tag Awareness in VPLS

4.1. No Duplicated MAC Address

One MAC address might be used by multiple hosts in different customer VLANs (C-VLAN). This is illegal but it is the headache reality for providers. In the virtualization environment, Virtual Machines (VM) are more likely to have duplicated MAC addresses. When these hosts/VMs join in the same VSI of a PE, the PE will see MAC address duplication. In order to overcome this issue, the PE has to adopt qualified learning [[RFC4762](#)], i.e., the PE has to set up one VSI per C-VLAN. This brings the scalability issue as discussed in [Section 4.2](#).

If the PE uses U-tag as the de-multiplexor to isolate traffic of customers' BDs, above MAC address duplication issue can be avoided.

4.2. Scalable Interconnection of L2 Sites

For the qualified learning, providers need set up one PW per C-VLAN. When there is a large number of customers multiplied by C-VLANs interconnected using VPLS, a mass of PWs need be maintained. It brings heavy operating overhead to providers.

In this document, U-tag is used to distinguish BDs in VPLS. In this way, traffic from multiple C-VLANs can be handled by a single VPLS. As shown in Figure 3.1, one PW is set up for each VSI and this VSI may be an integration of multiple BDs. Operating overhead of operators can be greatly reduced.

4.3. BUM Traffic Scoped per BD

Traditional VPLS limits a broadcast domain scope per PW. Suppose a customer has four sites in New York, Chicago, Atlanta and Dallas. BD1 = {New York, Chicago and Atlanta} while BD2 = {New York, Chicago and Dallas}. If one VSI per PE is set up to interconnect these four sites. BUM traffic of Atlanta site will be poured to Dallas site, and vice versa.

When PEs are aware of the U-tag, the BUM traffic can be confined per BD with multicast pruning. For above example, the operator need use two U-tags to distinguish the two BDs. In this way, BUM traffic of Atlanta site will be confined in BD1 and BUM traffic for Dallas site will be confined in BD2. This increases the efficiency of the bandwidth utilization of BUM traffic.

Two C-VLAN based multicast pruning techniques are listed below. (One is give in [[V-aware](#)] the other has been implemented by vendors.)

4.3.1. Advertising Interested VLANs in LDP

With the PW VLAN Vector TLV defined in [[V-aware](#)], PEs can advertise in LDP the interested C-VLANs for its interfaces. In this way, PEs can prune the flooding on a per C-VLAN basis.

4.3.2. Dynamic VLAN Registration with MVRP

It requires Multiple VLAN Registration Protocol (MVRP) to be supported by PEs for U-tag registration on the interfaces providing VPLS. With the help of MVRP, operators need not manually configure C-VLANs on PEs.

Only when a C-VLAN is registered in both directions of a PW, this PW will not be eliminated for this C-VLAN. Otherwise, this PW will be pruned for this C-VLAN. Multicast frames for a C-VLAN SHOULD only be forwarded on PWs that are not pruned for this C-VLAN.

4.4. Per C-VLAN MAC Withdraw

With the awareness of U-tag, PEs can achieve a finer grained C-VLAN scoped MAC withdraw. For example, with the VLAN Vector TLV defined in [\[V-aware\]](#), a PE can specify VLANs that it wants their MAC address to be flushed.

5. Backward Compatibility

Two PEs need negotiate their capability on supporting the awareness of U-tag. Unless both PEs are aware of U-tag, the tagged PW cannot be established. When a PE realizes the peering PE's interface is unaware of U-tag, it MUST fall back to establish a raw PW with this interface.

There are two ways to achieve the capability negotiation.

- a) As defined in Section 4 of [\[V-aware\]](#), PEs can negotiate this capability through LDP using the VLAN Aware Capability TLV.
- b) A tagged PW is established between two interfaces if they both enable MVRP.

For the tagged PW, PEs can achieve customer VLAN scoped MAC address flushing [\[V-aware\]](#). However, PEs may as well send out the old type MAC withdraw message per [Section 6.2 of \[RFC4762\]](#). The receiver PE parses this kind of message as that the peering PE is flushing MAC addresses across all customer VLANs supported by this PW.

6. Contributors

Xingjian He, Huawei

7. Security Considerations

This document raises no new security issues. For general security considerations, refer to [\[RFC4761\]](#) and [\[RFC4762\]](#).

8. IANA Considerations

This document requires no IANA actions. RFC Editor: please remove this section before publication.

9. References

9.1. Normative References

- [V-aware] D. Cai, S. Boutros, and et al, "VLAN Aware VPLS services", [draft-cai-l2vpn-vpls-vlan-aware-bundling-00.txt](#), working in progress.
- [MVRP] IEEE P802.1ak/D8.0, "IEEE Standard for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks -- Amendment 07: Multiple Registration Protocol", November 29, 2006.
- [RFC4762] Lasserre, M., Ed., and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC 4762](#), January 2007.

9.2. Informative References

- [RFC4761] Kompella, K., Ed., and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC 4761](#), January 2007.

Author's Addresses

Mingui Zhang
Huawei Technologies
No. 156 Beiqing Rd. Haidian District,
Beijing 100095
P.R. China

EMail: zhangmingui@huawei.com

Bin Wang
Huawei Technologies
No. 156 Beiqing Rd. Haidian District,
Beijing 100095
P.R. China

EMail: zhangmingui@huawei.com

Liang Xia
Huawei Technologies

Email: frank.xialiang@huawei.com

Jie Hu
China Telecom

Email: hujie@ctbri.com.cn

