

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 17, 2014

L. Zheng
Z. Li
S. Aldrin
Huawei Technologies
February 13, 2014

A Framework for E-VPN Performance Monitoring
draft-zheng-l2vpn-evpn-pm-framework-01

Abstract

The capability of Ethernet VPN performance monitoring (PM) is important to meet the Service Level Agreement (SLA) for the service beared. Since multipoint-to-point or multipoint-to-multipoint (MP2MP) network model applies, flow identifying is a big challenge for E-VPN PM. This document specifies the framework and mechanisms for the application of E-VPN PM.

Requirements Language

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](#)].

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 August 17, 2014.

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](http://trustee.ietf.org/license-info) 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	2
2.	Terminology	3
3.	Overview and Concepts	4
3.1.	EVI-to-EVI Tunnel	4
4.	Control Plane	4
4.1.	E-VPN Membership Auto-Discovery	4
4.2.	EVI-to-EVI Tunnel Label Allocation	4
5.	Data Plane	5
5.1.	Additional Label for Ingress EVI Identification	5
5.2.	Replace MAC Label with ET Label	6
6.	E-VPN Performance Monitoring	6
7.	IANA Considerations	7
8.	Security Considerations	7
9.	Acknowledgements	7
10.	References	7
10.1.	Normative References	7
10.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

Virtual Private LAN Service (VPLS) is a proven and widely deployed Ethernet L2VPN solution. However, it has a number of limitations when it comes to redundancy, multicast optimization and provisioning simplicity. Also, new applications are driving several new requirements for other L2VPN services such as E-TREE, VPWS, and VPMS. Furthermore, data center interconnect applications are driving the need for new service interface types, the "VLAN-aware Bundling" service interfaces. Then the Ethernet VPN (E-VPN) solution (defined in [[I-D.ietf-l2vpn-evpn](#)]) has been proposed to meet these requirements which is documented in [[I-D.ietf-l2vpn-evpn-req](#)].

An E-VPN comprises PEs that form the edge of the MPLS infrastructure and CEs that are connected to PEs. The PEs provide virtual Layer 2 bridged connectivity between the CEs. In E-VPN, MAC learning between PEs occurs not in the data plane but in the control plane. PEs

advertise the MAC addresses learned from the CEs, along with the associated MPLS label, to other PEs in the control plane by using MP-BGP.

The requirements and reference framework for Ethernet VPN (E-VPN) Operations, Administration and Maintenance (OAM) has been specified in [[I-D.salam-l2vpn-evpn-oam-req-frmwk](#)]. The capability of E-VPN to measure and monitor performance metrics for packet loss, packet delay, etc. is essential for meeting the Service Level Agreement (SLA). This measurement capability also provides operators with greater visibility into the performance characteristics of the services in their networks, and provides diagnostic information in case of performance degradation or failure and helps for fault localization. To perform the measurement of packet loss, delay and other metrics on a particular E-VPN flow, the egress PE needs to determine which specific ingress EVI packets belongs to. There exists complete and mature performance monitoring mechanism for the traditional L2VPN based on the point-to-point PW. But in the case of E-VPN, multipoint-to-point (MP2P) or multipoint-to-multipoint (MP2MP) network model applies, it makes the flow identifying a big challenge for packets loss and delay measurement. This MP2P or MP2MP model also apply to L3VPN, please refer to [[I-D.zheng-l3vpn-pm-analysis](#)] for detailed description of the challenge for performance monitoring of such network model.

Statistical approximation of packet loss by using synthetic OAM packets is briefly discussed in [[I-D.salam-l2vpn-evpn-oam-req-frmwk](#)]. This document defines a framework for accurate performance monitoring of E-VPN, especially for packet loss measurement. The point-to-point connection named as EVI-to-EVI tunnel is introduced in E-VPN. And the corresponding process of control plane and data plane is defined.

[2.](#) Terminology

E-VPN: Ethernet VPN

EVI: Ethernet VPN Instance

ET: EVI-to-EVI Tunnel

MP2P: Multi-Point to Point

MP2MP: Multi-Point to Multi-Point

P2P: Point to Point

PM: Performance Monitoring

[3.](#) Overview and Concepts

Based on the mechanisms in [[I-D.ietf-l2vpn-evpn](#)], for a particular MAC address route, the directly connected PE allocates the same MPLS label to all the remote PEs which maintain the MAC routing and forwarding instance (EVI) of that E-VPN. Thus for the egress PE, it is unable to identify the source EVI of the received E-VPN packets.

To perform the packet loss or delay measurement on a specific E-VPN flow, it is critical to establish the Point-to-Point connection between the two EVIs. Once the Point-to-Point connection is built up, current measurement mechanisms for MPLS networks may be applied to E-VPN. A new concept "EVI-to-EVI Tunnel" is introduced in the following section to establish such Point-to-Point connection in E-VPN.

[3.1.](#) EVI-to-EVI Tunnel

In order to perform performance monitoring in E-VPN, a point-to-point connection between any two EVIs of a particular E-VPN needs to be established. This point-to-point connection enables the egress PE identifying the ingress EVI of the received E-VPN packet, thus enables the measurement of the packet loss and delay between the ingress and egress EVIs. Such point-to-point connection between an ingress EVI and an egress EVI is called "EVI-to-EVI Tunnel (ET)".

[4.](#) Control Plane

This section describes the control plane mechanisms for E-VPN

performance monitoring.

[4.1.](#) E-VPN Membership Auto-Discovery

Before the Point-to-Point connections between EVIs could be established, each PE attaching a given E-VPN needs to learn all the remote PEs that attach to the same E-VPN. This could be achieved by the Ethernet A-D route per EVI defined in [[I-D.ietf-l2vpn-evpn](#)]. Please refer to [section 9.4.1](#) [[I-D.ietf-l2vpn-evpn](#)] for details.

[4.2.](#) EVI-to-EVI Tunnel Label Allocation

After obtaining the E-VPN membership information, each PE needs to allocate MPLS labels to identify the EVI-to-EVI tunnel from the remote EVI to the local EVI. We call such labels as ET labels in this document. For each local EVI, the egress PE SHOULD allocate different ET labels for each remote EVI in PEs belonging to the same E-VPN. As such, the egress PE could identify the E-VPN flow received from different ingress EVIs, and the packet loss and delay

Zheng, et al.

Expires August 17, 2014

[Page 4]

Internet-Draft

A Framework for EVPN PM

February 2014

measurement could be performed between each ingress EVI and the local EVI.

[5.](#) Data Plane

This section introduces two new MPLS label stack encapsulations when ET label applies.

[5.1.](#) Additional Label for Ingress EVI Identification

When a E-VPN data packet is to be sent on the ingress PE, firstly the label advertised by the MP-BGP for the Mac address route is pushed onto the label stack. The ET label allocated by the egress EVI for the ingress EVI should then be pushed onto the label stack to identify the Point-to-Point connection between the sending and receiving EVI. Finally the MPLS tunnel label is pushed onto the label stack. The process of TTL and COS fields between the E-VPN label encapsulation and the tunnel label encapsulation is done according to the Pipe and Uniform Models defined by [[RFC3270](#)] and [[RFC3443](#)]. The value of the TTL and COS field in the MAC label's encapsulation SHOULD be copied to the corresponding fields of the ET label's encapsulation. As such, one extra label is carried in the

label stack compared with E-VPN data plane defined in [\[I-D.ietf-l2vpn-evpn\]](#).

When the E-VPN data packet received by the egress PE, the outermost tunnel label is popped, then the egress PE could use the ET label to identify the ingress EVI of the packet. The process of TTL and COS fields at the egress node should be done according to the Pipe and Uniform Models defined by [\[RFC3270\]](#) and [\[RFC3443\]](#). Since the value of the TTL and COS fields of the MAC label encapsulation and the ET label encapsulation are the same, the TTL and COS fields of the ET label encapsulation could be ignored during the course of the TTL and COS process at the egress node.

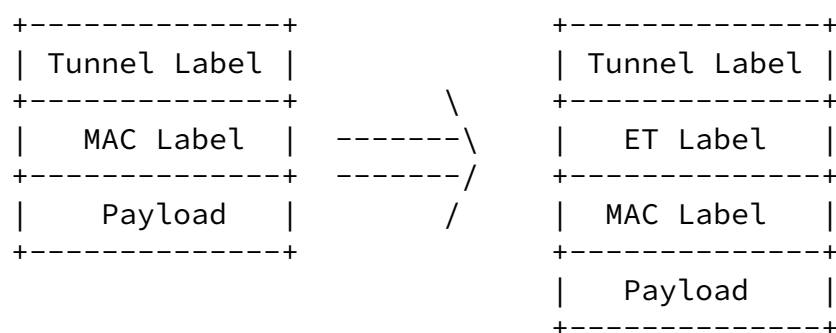


Fig.1 Additional Label for Ingress EVI Identification

[5.2.](#) Replace MAC Label with ET Label

Since the ET label identifies the connection between the ingress EVI and egress EVI, it could also be used to identify the egress EVI forwarding table in which the MAC prefix lookup should be performed. Thus when encapsulating the E-VPN data packets, the ingress PE could simply replace the MAC label with the ET label, then push the tunnel label. The process of TTL and COS fields between the MAC label encapsulation and the tunnel label encapsulation is done according to the Pipe and Uniform Models defined by [\[RFC3270\]](#) and [\[RFC3443\]](#). The TTL and COS value of the MAC label entry should be copied to the TTL and COS field of the ET label entry respectively. In this way the depth of the MPLS label stack is unchanged.

The encapsulation method would require the egress PE to perform MAC

prefix lookup in the egress EVI forwarding table before the packet can be forwarded to a specific CE. The similar procedure is also required when per-instance EVI label allocation mechanism is used. The process of TTL and COS fields at the egress node should be done according to the Pipe and Uniform Models defined by [RFC3270] and [RFC3443]. Since the MAC label encapsulation is replaced with the ET label encapsulation, the TTL and COS fields of the VT label encapsulation should be used as those of the MAC label encapsulation during the course of the TTL and COS process at the egress node.

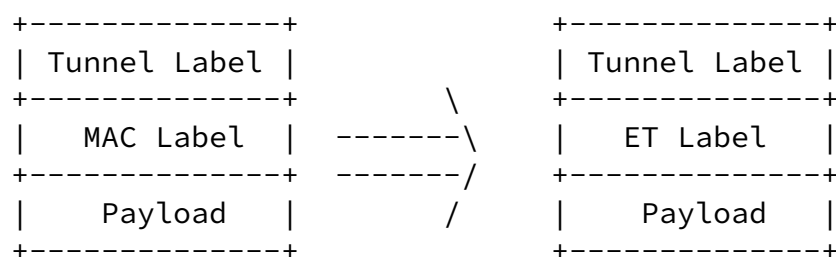


Fig.2 Replace the MAC Label with ET Label

6. E-VPN Performance Monitoring

[RFC6374] defines procedure and protocol mechanisms to enable the efficient and accurate measurement of packet loss, delay, as well as related metrics in MPLS networks. It provides either point-to-point or point-to-multipoint measurement capabilities. Once the point-to-point connection EVI-to-EVI Tunnel is established between the ingress and egress EVIs, the procedures for the packet loss and delay measurement as defined in [RFC6374] can be utilized for E-VPN performance monitoring. The main difference between performance monitoring of E-VPN and MPLS is the format of identifiers in the Loss Measurement (LM) and Delay Measurement (DM) messages. Specifically, for E-VPN, the source and destination addresses of the LM and DM

messages should be set to the concatenation of the Route Distinguisher (RD) of the particular EVI and the IP address of the ingress and egress PE respectively.

7. IANA Considerations

This document makes no request of IANA.

[8.](#) Security Considerations

TBD

[9.](#) Acknowledgements

TBD

[10.](#) References

[10.1.](#) Normative References

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

[10.2.](#) Informative References

[I-D.ietf-l2vpn-evpn-req]
Sajassi, A., Aggarwal, R., Bitar, N., and A. Isaac,
"Requirements for Ethernet VPN (EVPN)", [draft-ietf-l2vpn-evpn-req-07](#) (work in progress), February 2014.

[I-D.ietf-l2vpn-evpn]
Sajassi, A., Aggarwal, R., Henderickx, W., Isaac, A., and
J. Uttaro, "BGP MPLS Based Ethernet VPN", [draft-ietf-l2vpn-evpn-05](#) (work in progress), February 2014.

[I-D.salam-l2vpn-evpn-oam-req-frmwk]
Salam, S., Sajassi, A., Aldrin, S., and J. Drake, "E-VPN
Operations, Administration and Maintenance Requirements
and Framework", [draft-salam-l2vpn-evpn-oam-req-frmwk-02](#)
(work in progress), January 2014.

[I-D.zheng-l3vpn-pm-analysis]
Zheng, L., Li, Z., Aldrin, S., and B. Parise, "Performance
Monitoring Analysis for L3VPN", [draft-zheng-l3vpn-pm-analysis-02](#) (work in progress), October 2013.

P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", [RFC 3270](#), May 2002.

[RFC3443] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", [RFC 3443](#), January 2003.

[RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", [RFC 6374](#), September 2011.

Authors' Addresses

Lianshu Zheng
Huawei Technologies
Huawei Campus, No.156 Beiqing Rd.
Beijing 100095
China

Email: vero.zheng@huawei.com

Zhenbin Li
Huawei Technologies
Huawei Campus, No.156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Sam K. Aldrin
Huawei Technologies

Email: aldrin.ietf@gmail.com