Network Working Group J. Dong
Internet-Draft Z. Li
Intended status: Informational Huawei Technologies
Expires: April 18, 2013 October 15, 2012

A Framework for L3VPN Performance Monitoring
draft-dong-l3vpn-pm-framework-00

Abstract

This document specifies the framework and mechanisms for the application of performance monitoring (PM) to BGP/MPLS IP Virtual Private Networks (L3VPN).

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 April 18, 2013.

Copyright Notice

Table of Contents

1. Introduction

Level 3 Virtual Private Network (L3VPN) [RFC4364] service is widely deployed to provide enterprise VPN, Voice over IP (VoIP), video, mobile backhaul, etc. services. Most of these services are sensitive to the packet loss and delay. The capability to measure and monitor performance metrics for packet loss, delay, as well as related metrics 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 VPN traffic flow, the egress PE needs to identify the ingress VRF sending the VPN packets. As specified in the [L3VPN-PM-ANA] document, such flow identification is a big challenge for existing L3VPN.

This document specifies the framework and mechanisms for the application of performance monitoring in L3VPN.

2. Overview and Concepts

Based on the mechanisms in [RFC4364], for a particular VPN prefix, the directly connected PE allocates the same VPN label to all the remote PEs which maintain VPN Routing and Forwarding Tables (VRFs) of that VPN. Thus performance monitoring can not be performed on the egress PE, since it is not able to identify the source VRF of the received VPN packets.

As analyzed by [L3VPN-PM-ANA], to perform the packet loss or delay measurement on a specific VPN flow, it is critical for the egress PE to identify the unique VRF, i.e. to establish the Point-to-Point connection between the two VRFs . Once the Point-to-Point connection is built up, current measurement mechanisms may be applied to L3VPN. A new concept "VRF-to-VRF Tunnel" is introduced in the following section to establish such Point-to-Point connection.

2.1. VRF-to-VRF Tunnel

In order to perform performance monitoring in L3VPN, a point-to-point connection between any two VRFs of a particular VPN needs to be established. This guarantees that the egress PE could identify the ingress VRF of the received VPN traffic, thus it could measure the packet loss and delay between the ingress and egress VRFs. Such point-to-point VPN connection between an ingress VRF and an egress

VRF is called "VRF-to-VRF Tunnel (VT)".

3. Control Plane

This section describes the control plane mechanisms needed for L3VPN performance monitoring.

3.1. VPN Membership Auto-Discovery

Before establishing the Point-to-Point connections between VRFs, each PE needs to know all the remote PEs participating in the same VPN. This can be achieved by the membership auto-discovery procedure. Some mechanisms similar to the membership auto-discovery in VPLS [RFC4761] and L2VPN [RFC6074]can be used.

3.2. VRF-to-VRF Label Allocation

After obtaining the VPN membership information, each PE needs to allocate MPLS labels to identify the VRF-to-VRF tunnel between the local VRF and the remote VRFs, such labels are called VT labels. For each local VRF, the egress PE SHOULD allocate different VT labels for each remote VRF in PEs belonging to the same VPN. This way, the egress PE could identify the VPN flow received from different ingress VRFs, and the packet loss and delay measurement could be performed between each ingress VRF and the local VRF.

4. Data Plane

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

4.1. Additional Label for Ingress VRF Identification

When a VPN data packet needs to be sent, firstly the VPN label obtained from the BGP VPN route of the destination address prefix is pushed onto the label stack. The VT label allocated by the egress VRF should then be pushed onto the label stack to identify the Point-to-Point connection between the sending and receiving VRF. Lastly, the MPLS tunnel label is pushed onto the label stack. The TTL and COS value in the VPN label entry should be copied to the TTL and COS fields of the VT label respectively. This way, one additional label is carried in the label stack compared with L3VPN data plane in [RFC4364].

When the VPN data packet arrives at the egress PE, the outermost tunnel label is popped, then the egress PE could use the VT label to

identify the ingress VRF of the packet.

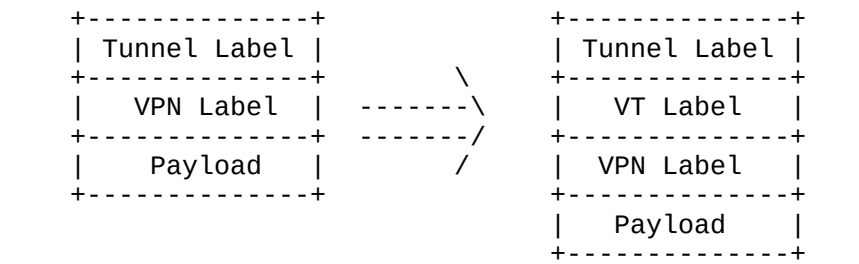

Fig.1 Additional Label for Ingress VRF Identification

4.2. Replace the VPN Label with VT Label

Since the VT label identifies the connection between the ingress VRF and egress VRF, it could also be used to identify the egress VRF table in which the VPN prefix lookup should be performed. Thus when encapsulating the VPN data packets, the ingress PE could simply replace the VPN label with the VT label, then push the tunnel label. The TTL and COS value of the VPN label entry should be copied to the TTL and COS field of the VT label respectively. This way the depth of the MPLS label stack is unchanged. Though this would require the egress PE to perform VPN prefix lookup in the egress VRF table before the packet can be forwarded to a specific CE, such lookup procedure is also required when per-instance VPN label allocation mechanism is used.

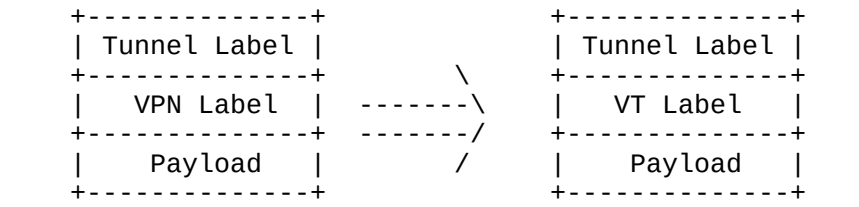

Fig.2 Replace the VPN Label with VT Label

5. L3VPN Performance Monitoring

Since the challenge of identifying the ingress VRF is resolved in section 4, the procedures for the packet loss and delay measurement as defined in [RFC6374] can be utilized for L3VPN performance monitoring. The main difference between performance monitoring of L3VPN and MPLS is the format of identifiers in the Loss Measurement (LM) and Delay Measurement (DM) messages. Specifically, for L3VPN, 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 VRF and the IP address of the ingress and egress PE respectively.

6. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC.

7. Security Considerations

TBD

8. Acknowledgements

TBD

9. References

9.1. Normative References

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

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.

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

9.2. Informative References

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

[RFC6074] Rosen, E., Davie, B., Radoaca, V., and W. Luo, "Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)", RFC 6074, January 2011.

Authors' Addresses

Jie Dong
Huawei Technologies
Huawei Building, No.156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

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

Email: lizhenbin@huawei.com