Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 21, 2016

J. Dong
Z. Li
Huawei Technologies
B. Parise
Cisco Systems
October 19, 2015

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

Abstract

The performance monitoring (PM) of BGP/MPLS IP Virtual Private Networks (L3VPN) is important for satisfying the Service Level Agreement(SLA) for critical network services. Since L3VPN is essentially using a multipoint-to-point service model, flow identification becomes a big challenge for L3VPN PM. This document specifies the framework and mechanisms for the application of PM in 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 21, 2016.

Table of Contents

1. Introduction

BGP/MPLS IP Virtual Private Networks (L3VPN) [RFC4364] is widely deployed to provide various services such as enterprise VPN, Voice over IP (VoIP), video, mobile backhaul, etc. Most of these services are sensitive to packet loss and delay. The capability to measure and monitor the performance metrics such as packet loss, delay, as well as related metrics is important for meeting the Service Level Agreements (SLA). This performance 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 fault localization.

In order to perform the measurement of packet loss, delay and other metrics on a particular L3VPN flow, the egress PE needs to identify the ingress VRF sending the VPN packets. As specified in [I-D.zheng-l3vpn-pm-analysis], such flow identification is a big challenge for L3VPN.

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

2. Flow Identification in L3VPN PM

Based on the mechanisms defined in [RFC4364], for a specific VPN prefix, the directly connected PE would allocate and advertise the same VPN label to all the remote PEs which have the VPN Routing and Forwarding Tables (VRFs) of the same VPN. Essentially this is a multipoint-to-point service model. On the egress PE, performance monitoring can not be performed based on the VPN label, because it cannot identify the ingress VRF which generates the VPN packets.

As analyzed in [I-D.zheng-l3vpn-pm-analysis], in order to perform the packet loss or delay measurement on a specific L3VPN traffic flow, it is critical that the egress PE can uniquely identify the ingress VRF of the received VPN packets.

[I-D.bryant-mpls-synonymous-flow-labels] defines the concept of Synonymous Flow Labels (SFL), for which the typical use case is the performance monitoring of MPLS applications. The Synonymous Flow Labels are used here for the performance monitoring of L3VPN, in which the SFLs are used by the egress PE to uniquely identify the ingress VRF of the received VPN packets. Depends on specific provisioning models, the SFLs can be either local-allocated or globally allocated labels. Subsequent sections specifies the data plane encapsulation and control plane considerations for different modes.

3. Local-Allocated SFL for L3VPN PM

This section specifies the L3VPN PM mechanism with local-allocated Synonymous Flow Labels, in which the SFLs are allocated by the egress PEs. The SFL can be allocated to identify a specific ingress VRF, or a specific ingress-egress VRF pair. Depends on the semantics of the SFL, two MPLS label stack encapsulations are used.

3.1. Additional SFL for Source Identification

This section specifies the label stack encapsulation in which the SFL is allocated by the egress PE to identify the ingress VRF and the ingress PE. While the VPN label is still used for packet forwarding decision on the egress PE.

When a VPN data packet is to be sent by an ingress PE, firstly the VPN label obtained from the BGP VPN route of the destination address prefix is pushed onto the label stack. Then according to the next-hop of the BGP VPN route, the Synonymous Flow Label allocated by the next-hop PE for the ingress VRF SHOULD be pushed onto the label stack. Finally, the MPLS tunnel label is pushed onto the label stack. The TTL and TC fields of the VPN label and the tunnel label entries SHOULD be set according to the Pipe or Uniform Model as defined in [RFC3270] and [RFC3443]. The value of the TTL and TC fields of the VPN label entry SHOULD be copied to the TTL and TC fields of the Synonymous Flow Label entry respectively. With this encapsulation, one additional label is carried in the label stack compared with traditional L3VPN encapsulation defined in [RFC4364].

When the VPN packet arrives at the egress PE, the outermost tunnel label is popped (if present), then the egress PE uses the Synonymous Flow Label to identify the ingress VRF of the packet. The TTL and COS fields SHOULD be processed according to the Pipe or Uniform Models defined in [RFC3270] and [RFC3443]. Since the value of the TTL and TC fields of the VPN label and the SFL are the same, the TTL and TC fields of the SFL can be ignored by the egress PE.

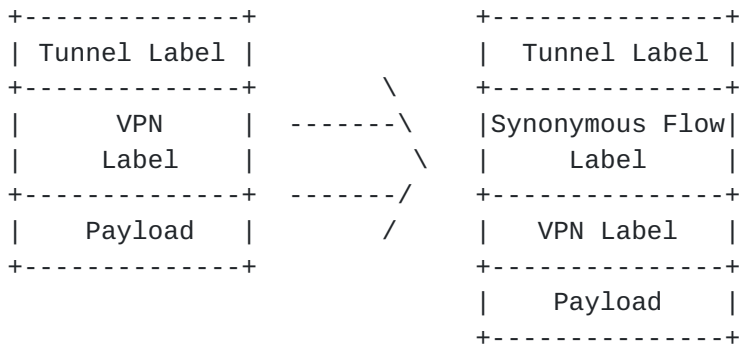

Figure 1. Additional SFL for Source Identification

3.2. Replacing the VPN Label with SFL

This section specifies the label stack encapsulation in which the SFL is allocated by the egress PE to identify a specific ingress-egress VRF pair. In this case, since the SFL could also be used by the egress PE to identify the egress VRF, if the VPN label is a per-

instance label, on the ingress PE the VPN label can be replaced with the SFL, then the tunnel label is pushed onto the label stack. The TTL and TC fields of the Synonymous Flow Label and the tunnel label SHOULD be set according to the Pipe or Uniform Model as defined in [RFC3270] and [RFC3443]. The value of the TTL and TC fields of the VPN label entry should be copied to the TTL and TC fields of the Synonymous Flow Label entry respectively. With this mechanism, the depth of the MPLS label stack is not increased, while the number of Synonymous Flow Labels needed would be more than that in the mechanism of Section 3.1.

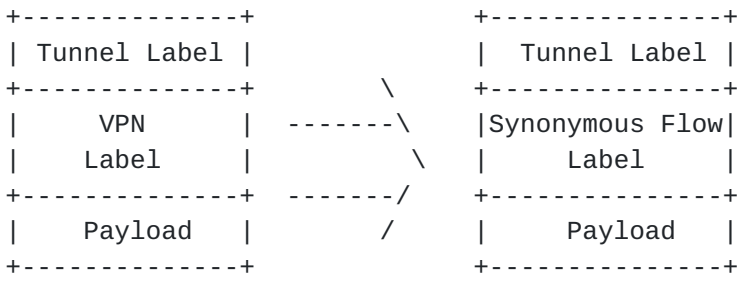

Figure 2. VPN label replaced with SFL

Note that this label stack encapsulation would require the egress PE to lookup the destination VPN prefix in the egress VRF before the packet can be forwarded to a specific CE. This is similar to the per-instance VPN label allocation mechanism in [RFC4364]. The TTL and TC fields SHOULD be processed according to the Pipe or Uniform Model as defined in [RFC3270] and [RFC3443]. Since the VPN label entry is replaced with the Synonymous Flow Label, the TTL and TC fields of the SFL should be used as those of the VPN label entry in traditional L3VPN encapsulation.

4. Global SFL for L3VPN PM

In some scenarios global MPLS label can be beneficial for L3VPN services. This section specifies the L3VPN PM mechanism with global Synonymous Flow Labels.

4.1. Global SFL for VPN Identification

In this mode, a global unique SFL is allocated for each VPN. Besides, a global unique label is allocated to identify each PE node in the network. An ingress VRF can be identified by the combination of the PE label and the VPN SFL label.

When a VPN data packet is to be sent by an ingress PE, firstly the VPN label obtained from the BGP VPN route of the destination address prefix is pushed onto the label stack. Then the Synonymous VPN Label

and the PE label are pushed onto the label stack. Finally, the MPLS tunnel label is pushed onto the label stack.

With this approach, two additional labels are carried in the label stack compared with traditional L3VPN encapsulation defined in [RFC4364].

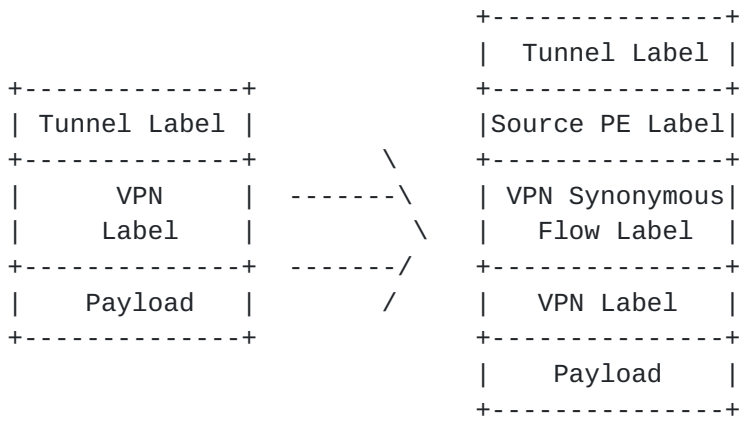

Figure 3. Label stack with global PE label and VPN SFL

In scenarios where the VPN label is per-instance label, the VPN Synonymous Flow Label can replace the VPN label, then the MPLS label stack would contain one additional label compared with traditional L3VPN encapsulation.

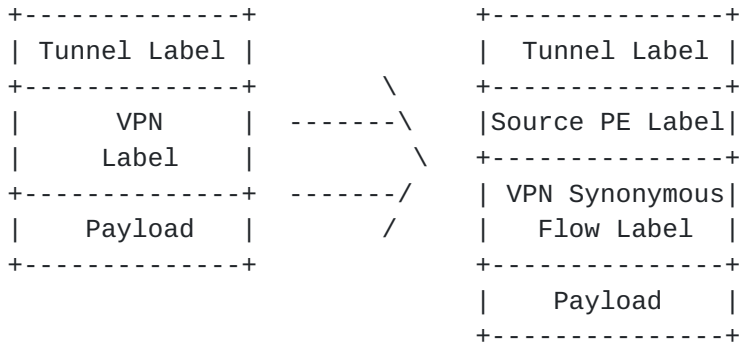

Figure 4. VPN label replaced with global VPN SFL

4.2. Global SFL for Ingress VRF Identification

In this mode, a global unique SFL is allocated for each VRF on each PE, which can be used by the egress PE to identify both the ingress VRF and the ingress PE.

When a VPN data packet is to be sent by an ingress PE, firstly the VPN label obtained from the BGP VPN route of the destination address

prefix is pushed onto the label stack. Then the VRF SFL is pushed onto the label stack. Finally, the MPLS tunnel label is pushed. With this approach, one additional label is carried in the label stack compared with the traditional L3VPN encapsulation.

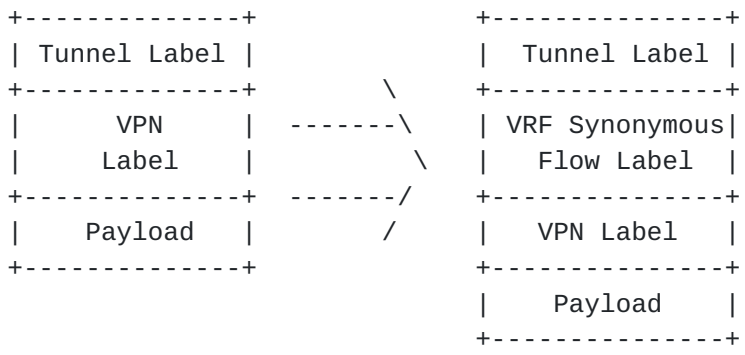

Figure 5. Label stack with global VRF SFL

When the VPN packet arrives at the egress PE, the outermost tunnel label is popped (if present), then the egress PE uses the VRF Synonymous Flow Label to identify the ingress VRF of the packet. The VPN label is used for packet forwarding decision on the egress PE.

5. Control Plane

This section describes the corresponding control plane mechanisms for L3VPN performance monitoring with the help of Synonymous Flow Label.

5.1. VPN Membership Auto-Discovery

Before the Synonymous Flow Labels are allocated, a PE which attaches to a particular VPN needs to know all the remote VRFs on other PEs that attach to the same VPN. This is achieved via the BGP membership auto-discovery procedure. Mechanisms similar to the membership auto-discovery of MVPN [RFC6513] can be used. Detailed BGP protocol extensions will be specified in a companion document.

5.2. Allocation of Synonymous Flow Label for L3VPN PM

After the VPN membership information is obtained, Synonymous Flow Labels needs to be allocated for L3VPN PM. The SFL can be either local-allocated by each PE, or it can be global unique label which is allocated for L3VPN PM.

5.2.1. Local SFL Allocation

With local label allocation, for each attached VPN, a PE SHOULD allocate unique Synonymous Flow Label for each remote VRF on its remote PEs. Two label allocation methods can be used:

1. A Synonymous Flow Label is allocated for each remote VRF. This SFL would be used by the egress PE to identify the ingress VRF of the received packet.

2. A Synonymous Flow Label is allocated for each remote-local VRF pair. With this approach, the SFL can replace the VPN label in the MPLS label stack of L3VPN packet as specified in Section 3.2.

With both methods, the allocated Synonymous Flow Label SHOULD be advertised to remote PEs via the L3VPN control plane, where some extensions to BGP is needed. Detailed BGP protocol extensions will be specified in a future version.

5.2.2. Global SFL Allocation

With global label allocation, the SFLs are allocated by a network controller, which obtains the VPN membership information via the VPN membership auto-discovery. Two global label allocation methods can be used:

1. A global SFL is allocated for each VPN. The combination of this SFL and the global PE label can identify the ingress VRF of the received packets.

2. A global SFL is allocated for each VRF on each PE. This SFL itself can identify the ingress VRF of the received packets.

Detailed mechanisms about the global SFL allocation will be specified in a companion document.

6. L3VPN Performance Monitoring

Since the challenge of source identification in L3VPN is resolved, the procedures for the packet loss and delay measurement as defined in [RFC6374] can be applied to L3VPN performance monitoring. Note that in L3VPN performance monitoring, the source and destination address TLV of the LM and DM messages SHOULD be set to the VPN-IPv4 or VPN-IPv6 address, which begins with the 8-byte Route Distinguisher (RD) of the VRF and ends with a 4-byte IPv4 address or 16-byte IPv6 address of the ingress or egress PE node.

7. IANA Considerations

This document makes no request of IANA.

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

8. Security Considerations

TBD

9. References

9.1. Normative References

[I-D.bryant-mpls-synonymous-flow-labels]
Bryant, S., Swallow, G., Sivabalan, S., Mirsky, G., Chen, M., and Z. Li, "RFC6374 Synonymous Flow Labels", draft-bryant-mpls-synonymous-flow-labels-01 (work in progress), July 2015.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC3270] Le Faucheur, F., Wu, L., Davie, B., Davari, S., Vaananen, P., Krishnan, R., Cheval, P., and J. Heinanen, "Multi-Protocol Label Switching (MPLS) Support of Differentiated Services", RFC 3270, DOI 10.17487/RFC3270, May 2002, <http://www.rfc-editor.org/info/rfc3270>.

[RFC3443] Agarwal, P. and B. Akyol, "Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks", RFC 3443, DOI 10.17487/RFC3443, January 2003, <http://www.rfc-editor.org/info/rfc3443>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <http://www.rfc-editor.org/info/rfc4364>.

[RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, DOI 10.17487/RFC6374, September 2011, <http://www.rfc-editor.org/info/rfc6374>.

9.2. Informative References

[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-03 (work in progress), July 2014.

[RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <http://www.rfc-editor.org/info/rfc6513>.

Authors' Addresses

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

Email: jie.dong@huawei.com

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

Email: lizhenbin@huawei.com

Bhavani Parise
Cisco Systems

Email: bhavani@cisco.com