

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 18, 2013

L. Zheng
Z. Li
Huawei Technologies
October 15, 2012

Performance Monitoring Analysis for L3VPN
draft-zheng-l3vpn-pm-analysis-00

Abstract

To perform the measurement of packet loss, delay and other metrics on a particular VPN flow, the egress PE need to tell to which specific ingress VRF a packet belongs. But for L3VPN, multipoint-to-point or multipoint-to-multipoint (MP2MP) network model applies, flow identifying is a big challenge. This document summarizes the current performance monitoring mechanisms for MPLS networks, and analyzes the challenge for L3VPN performance monitoring. This document also discuss the key points need to be taken in consideration when designing L3VPN performance monitoring mechanisms.

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

Copyright (c) 2012 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	4
2.	Overview of current mechanisms for MPLS networks	5
2.1.	Packet Loss and Delay Measurement for MPLS Networks	5
2.2.	Profile for MPLS-based Transport Networks	5
3.	Challenge for L3VPN Performance Monitoring	6
4.	Design Consideration	8
4.1.	P2P Connection	8
4.2.	Control Plane	8
4.3.	Data Plane	8
4.4.	MPLS OAM	8
4.5.	QoS	9
4.6.	Configuration	9
5.	Security Considerations	10
6.	IANA Considerations	11
7.	Acknowledgements	12
8.	References	13
8.1.	Normative References	13
8.2.	Informative References	13
	Authors' Addresses	14

1. Introduction

Level 3 Virtual Private Network (L3VPN) [RFC4364] service is widely deployed in the production network. It is deployed to provide enterprise interconnection, Voice over IP (VoIP), video, mobile, 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 SLA. The requirement for SLA measurement for MPLS networks has been documented in [RFC4377].

One popular deployment of L3VPN nowadays is in mobile backhaul networks. When deploying MPLS-TP in mobile backhaul network, due to the scalability issue with PW, L3VPN is used either for end-to-end service delivery, or L2VPN and L3VPN hybrid networking. The measurement capability of L3VPN provides operators with greater visibility into the performance characteristics of 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 flow, the egress PE need to tell to which specific ingress VRF a packet belongs. But for L3VPN, multipoint-to-point or multipoint-to-multipoint (MP2MP) network model applies, flow identifying is a big challenge. This document summarizes the current performance monitoring mechanisms for MPLS networks, and analyzes the challenge for L3VPN performance monitoring. This document also discuss the key points need to be taken in consideration when designing L3VPN performance monitoring mechanisms.

2. Overview of current mechanisms for MPLS networks

2.1. Packet Loss and Delay Measurement for MPLS Networks

[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.

The LM protocol can perform two distinct kinds of loss measurement. In inferred mode, it can measure the loss of specially generated test packets (in order to infer the approximate data-plane loss level). In direct mode, it can directly measure data-plane packet loss. Direct measurement provides perfect loss accounting, but may require specialized hardware support and is only applicable to some LSP types. Inferred measurement provides only approximate loss accounting but is generally applicable.

The LM and DM protocols are initiated from a single node, the querier. A query message may be received either by a single node or by multiple nodes; i.e. these protocols provide point-to-point or point-to-multipoint measurement capabilities.

2.2. Profile for MPLS-based Transport Networks

Procedures for the measurement of packet loss, delay, and throughput in MPLS networks are defined in [RFC6374]. [RFC6375] describes a profile, i.e. a simplified subset, of procedures that suffices to meet the specific requirements of MPLS-based transport networks [RFC5921] as defined in [RFC5860]. This profile is presented for the convenience of implementors who are concerned exclusively with the transport network context.

LM session is externally configured and the values of several protocol parameters can be fixed in advance at the endpoints involved in the session, so that inspection or negotiation of these parameters is not required.

3. Challenge for L3VPN Performance Monitoring

To perform the measurement of packet loss, delay and other metrics on a particular VPN flow, the egress PE need to tell to which specific ingress VRF a packet belongs.

The above mentioned existing mechanisms for MPLS networks provide either point-to-point or point-to-multipoint measurement capabilities. For a specific receiver, it could easily identify a specific flow by the label stack information, when Penultimate Hop Pop (PHP) function is disabled .

But in the case of L3VPN, multipoint-to-point or multipoint-to-multipoint (MP2MP) network model applies , it makes the flow identifying a big challenge for packets loss and delay measurement. According to the label allocation mechanisms of L3VPN, a private label itself cannot uniquely identify a specific VPN flow. That is, when the egress PE allocates VPN label for a specific prefix of a VPN, the same label will be advertised to all its peers. Given a VPN flow, the egress PE cannot tell which ingress VRF is from based on the private label it carries. As a result, it's not feasible to perform the loss or delay measurement on this flow.

In L3VPN the LSPs may be merged at any intermediate nodes along the LSP (e.g., Label Distribution Protocol (LDP) [RFC5036] based LSP). The egress PE cannot derive a unique identifier of the source PE from label stack. The tunnel label cannot help for flow identification due to the LSP merge.

In L3VPN, the ingress PE could be identified by the tunnel label when TE LSP applies [RFC3209], but the egress PE cannot tell to which specific VRF a packet belongs when extranet (If the various sites in a VPN are owned by different enterprises) exist on ingress PE. Figure 1 shows an example of extranet. In Figure1, Site A,B,C,D all belong to the same VPN-A, but Site C and Site D does not belong to the same enterprise (Site D also belongs to a VPN-B), so different VRFs are maintained for each site on PE3. PE1 assign the same label L for prefix 10.0.0.1 to PE3 of VPN-A, when it recieve the VPN-A flow from PE3, it can not tell the flow is from either VRFC or VRFD by the label stack.

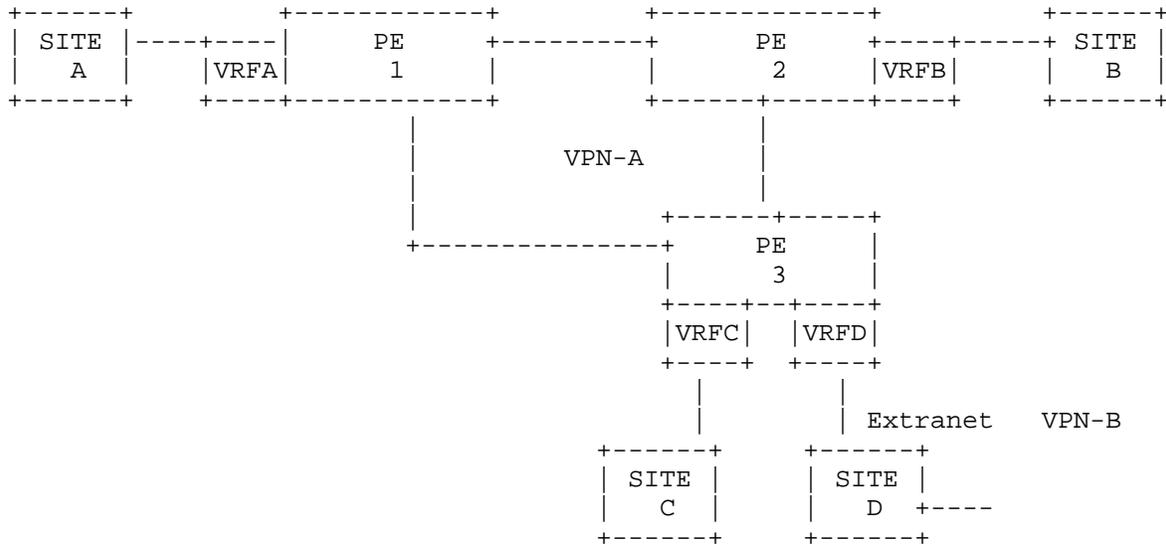


Figure1: Extranet on Ingress PE

The current label allocation mechanism of L3VPN make the flow identification a big challenge for L3VPN performace monitoring, as a result the current performace monitoring mechanisms for MPLS networks cannot be applied to L3VPN networks. Extension or alteration to current label allocation mechanism is needed to solve the problem.

4. Design Consideration

This section discuss the key points need to be taken in consideration when designing L3VPN performance monitoring mechanism.

4.1. P2P Connection

As analyzed above, to perform the packet loss or delay measurement on a specific VPN flow, it is critical for the egress PE to identify the unique ingress VRF, i.e. to establish the Point-to-Point connection between the two VRFs. Current allocation mechanism may need extension or alteration to help build up the Point-to-Point connection. Once the Point-to-Point connection is built up, current measurement mechanisms may be applied to L3VPN .

Conditions like Penultimate Hop Popping (PHP), Equal-Cost Multi-Path (ECMP) load-balancing and BGP multi-path may make it infeasible for receiving PE to identify the ingress PE. These conditions SHOULD be excluded for consideration for mechanism design.

4.2. Control Plane

In L3VPN, BGP is used to distribute a particular route, as well as an MPLS label that is mapped to that route [RFC4364]. The label mapping information for a particular route is piggybacked in the same BGP Update message that is used to distribute the route itself. In order to setup the Point-to-Point connection between ingress and egress VRFs the current label distribution mechanism may be altered. For compatibility, this alteration SHOULD NOT change the current label distribution mechanism dramatically.

4.3. Data Plane

Same as for control plane, for compatibility reason, the data plane should as far as be compatible with the current L3VPN forwarding procedure.

4.4. MPLS OAM

[RFC6374], [RFC6375] defines procedure and protocol mechanisms to enable the measurement of packet loss, delay, as well as related metrics in MPLS networks. These mechanisms SHOULD be reasonably reused in L3VPN networks. The addressing of source and destination of Loss Measurement (LM) and Delay Measurement (DM) messages may needed to be changed to identify the measured VRF.

4.5. QoS

To perform the packet loss or delay measurement in L3VPN network, either proactive or on-demand, SHOULD NOT impact the customer QoS experience.

4.6. Configuration

Measurement entities and functions MUST be configurable either statically or dynamically. It SHOULD be possible to configure and activated/deactivated the measurement capability as part of connectivity establishment, and it SHOULD also be possible to configure and activated/deactivated the capability after connectivity has been established .

5. Security Considerations

This document does not change the security properties of L3VPN.

6. IANA Considerations

This document makes no request to IANA.

7. Acknowledgements

The authors would like to thank XXX for their valuable comments.

8. References

8.1. Normative References

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

8.2. Informative References

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.
- [RFC4377] Nadeau, T., Morrow, M., Swallow, G., Allan, D., and S. Matsushima, "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC 4377, February 2006.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", RFC 5036, October 2007.
- [RFC5860] Vigoureux, M., Ward, D., and M. Betts, "Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks", RFC 5860, May 2010.
- [RFC5921] Bocci, M., Bryant, S., Frost, D., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, July 2010.
- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011.
- [RFC6375] Frost, D. and S. Bryant, "A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks", RFC 6375, September 2011.

Authors' Addresses

Lianshu Zheng
Huawei Technologies
China

Email: vero.zheng@huawei.com

Zhenbin Li
Huawei Technologies
China

Email: lizhenbin@huawei.com

