

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 3, 2015

M. Chen
Huawei
S. Ning
Tata Communications
V. Hallivuori
Tellabs Sinimaentie
December 30, 2014

BFD for MPLS LSPs Enhancement
draft-chen-mpls-bfd-enhancement-03

Abstract

This document defines extensions to BFD for MPLS LSPs that can be used to specify the return path of BFD control packets for a specific BFD session. Specifically, it re-uses the Reply Path TLV defined in [RFC7110](#) to specify the return path. Specifying congruent or more reliable return path increases robustness of BFD failure detection and reduces false positives.

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 July 3, 2015.

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](#) 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.	Session Establishment	3
3.	Using of Tunnel sub-TLVs	3
4.	IANA Considerations	4
5.	Security Considerations	4
6.	Acknowledgements	4
7.	Contributors	4
8.	References	4
8.1.	Normative References	4
8.2.	Informative References	5
	Authors' Addresses	5

[1.](#) Introduction

Bi-directional Forwarding Detection (BFD) [[RFC5880](#)] for Multiprotocol Label Switching (MPLS) is defined in [[RFC5884](#)], where Label Switched Path (LSP) Ping [[RFC4379](#)] is used to bootstrap the BFD session.

In traffic engineered networks, LSPs usually do not follow IP routing or there are multiple LSPs between same nodes. There is requirement to specify return path for a BFD session. For example, selecting a more reliable return path (e.g., TE LSP with Fast ReRoute (FRR) protection) could improve the reliability of reply BFD control packets or selecting co-routed LSP would ensure that failures in other return paths (e.g. shortest path for IP routing) do not cause BFD failure detection for the forward path.

This document defines extensions to BFD for MPLS LSPs that can be used to specify the return path of BFD control packets for a specific BFD session.

2. Session Establishment

All the procedures for session establishment defined in [\[RFC5884\]](#) apply here. In addition to that, a RP TLV defined in [\[RFC7110\]](#) MUST be carried in the echo request and the reply mode MUST be set to "Reply via specified path". This is used to notify the egress LSR to select the desired return path of the BFD session.

The specific return path is identified by the FEC sub-TLV encoded in the RP TLV. When the return path is required to follow the reverse direction of a bidirectional LSP (co-routed or associated). It just needs to set the B bit of RP TLV, no specific reply path FEC sub-TLV is carried.

When received an echo request with the reply mode set to "Reply via specified path" [\[RFC7110\]](#), if the receiver does not know the reply mode, an echo reply with the return code set to "Malformed echo request" and the Subcode set to zero SHOULD be send back to the initiator. The BFD session will not be established in this case.

When provisioning a single BFD to a bidirectional LSP, in case of the forwarding direction of session is OK but the reverse direction is not, the initiator MUST not send BFD control packets on the session until the echo reply is received from the terminator. And the terminator MUST not send BFD control packets onto the BFD session until the first BFD control packet is received from the initiator.

3. Using of Tunnel sub-TLVs

IPv4/IPv6 RSVP and Static Tunnel sub-TLVs are defined in [\[RFC7110\]](#), which are used to notify the egress LSR of a specific LSP how to choose the return path for an echo reply.

In this document, these sub-TLVs are re-used to notify the egress LSR of a specific LSP how to choose the return path of a BFD session.

As stated above, it is possible to have multiple parallel LSPs between the ingress and egress LSRs. In a typical network, a group of parallel LSPs have the same tunnel ID but different LSP IDs. There is one primary and one or more secondary LSPs in each direction. When establishing a BFD session for an LSP, it is common for the primary LSP to choose the primary as the return path and the secondary LSP to choose the secondary as the return path. With the IPv4/IPv6 Tunnel sub-TLVs, operators don't have to specify a specific return LSP. They just need to specify from which tunnel the return LSP should be chosen, and specify the type (e.g., primary or secondary) of the return LSP. The egress LSR will then choose the proper return path.

Use of IPv4/IPv6 RSVP and Static Tunnel sub-TLVs permits Make-Before-Break (MBB) on return path of BFD session. If a specific LSP other than Tunnel would be selected, return path would be interrupted if egress node would do MBB for the return path. This is due to the fact that MBB is based on signaling new LSP with different LSP ID and then deleting old LSP - after MBB selected return path would no longer be available.

4. IANA Considerations

This document makes no request of IANA.

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

5. Security Considerations

Security considerations discussed in [[RFC5884](#)], [[RFC5880](#)], [[RFC5883](#)], [[RFC4379](#)] and [[RFC7110](#)] apply to this document.

6. Acknowledgements

The authors would like to thank Greg Mirsky, Xinchun Guo and Wei Cao for their review and comments to this document.

7. Contributors

Vishwas Manral

IP Infusion

Almora, Uttarakhand India

EMail: vishwas@ipinfusion.com

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.
- [RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", [RFC 5884](#), June 2010.

- [RFC7110] Chen, M., Cao, W., Ning, S., Jounay, F., and S. Delord, "Return Path Specified Label Switched Path (LSP) Ping", [RFC 7110](#), January 2014.

8.2. Informative References

- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", [RFC 4379](#), February 2006.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", [RFC 5880](#), June 2010.
- [RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for Multihop Paths", [RFC 5883](#), June 2010.

Authors' Addresses

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

So Ning
Tata Communications

Email: ning.so@tatacommunications.com

Ville Hallivuori
Tellabs Sinimaentie

Email: ville.hallivuori@tellabs.com

