SPRING Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 9, 2017

G. Mirsky
ZTE Corp.

May 8, 2017

Bidirectional Forwarding Detection (BFD) in Segment Routing Networks Using MPLS Dataplane
draft-mirsky-spring-bfd-00

Abstract

Segment Routing architecture leverages the paradigm of source routing. It can be realized in the Multiprotocol Label Switching (MPLS) network without any change to the data plane. A segment is encoded as an MPLS label and an ordered list of segments is encoded as a stack of labels. Bidirectional Forwarding Detection (BFD) is expected to monitor any kind of paths between systems. This document defines how to use Label Switched Path Ping to bootstrap and control path in reverse direction of a BFD session on the Segment Routing network over MPLS dataplane.

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 November 9, 2017.

Copyright Notice

Table of Contents

1. Introduction

[RFC5880], [RFC5881], and [RFC5883] established the Bidirectional Forwarding Detection (BFD) protocol for IP networks. [RFC5884] and [RFC7726] set rules of using BFD Asynchronous mode over Multiprotocol Label Switching (MPLS) Label Switched Path (LSP). These latter standards implicitly assume that the egress BFD peer, which is the egress Label Edge Router (LER), will use the shortest path route regardless of the path the ingress LER uses to send BFD control packets towards it.

This document defines use of LSP Ping for Segment Routing networks over MPLS dataplane [I-D.ietf-mpls-spring-lsp-ping] to bootstrap and control path of a BFD session from the egress to ingress LER.

1.1. Conventions used in this document

1.1.1. Terminology

BFD: Bidirectional Forwarding Detection

FEC: Forwarding Equivalence Class

MPLS: Multiprotocol Label Switching

LSP: Label Switching Path

LER: Label Edge Router

1.1.2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Bootstrapping BFD session over Segment Routed tunnel

As discussed in [I-D.ietf-mpls-spring-lsp-ping] introduction of Segment Routing network domains with an MPLS data plane adds three new sub-TLVs that MAY be used with Target Forwarding Equivalence Class (FEC) TLV. Section 6.1 addresses use of the new sub-TLVs in Target FEC TLV in LSP ping and LSP traceroute. For the case of LSP ping the [I-D.ietf-mpls-spring-lsp-ping] states that:

> Initiator MUST include FEC(s) corresponding to the destination segment.
>
> Initiator, i.e. ingress LSR, MAY include FECs corresponding to some or all of segments imposed in the label stack by the ingress LSR to communicate the segments traversed.

It has been noted in [RFC5884] that a BFD session monitors for defects particular <MPLS LSP, FEC> tuple. [RFC7726] clarified how to establish and operate mutiple BFD sessions for the same <MPLS LSP, FEC> tuple. Because only ingress edge router is aware of the SR-based explicit route egress edge router can associate the LSP ping with BFD Discriminator TLV with only one of the FECs it advertised for the particular segment. Thus this document defines that: When LSP ping is used to bootstrap a BFD session this document updates the statement and defines that:

> When LSP Ping is used to bootstrap a BFD session it MUST include only one FEC corresponding to the destination segment and SHOULD NOT include FECs corresponding to some or all of other segments

imposed by the ingress LSR.

Operationally such restriction would not cause any problem or uncertainty as LSP ping with FECs corresponding to some or all segments or traceroute that validate the segment route MAY precede the LSP ping that bootstraps the BFD session.

Encapsulation of a BFD Control packet in Segment Routing network with MPLS dataplane MUST follow Section 7 [RFC5884] when IP/UDP header used and MUST follow Section 3.4 [RFC6428] without IP/UDP header being used.

3. Use BFD Reverse Path TLV over Segment Routed MPLS tunnel

When a BFD session is used to monitor a source routed unidirectional path there may be a need to direct egress BFD peer to use specific path for the reverse direction of the BFD session by using the BFD Reverse Path TLV [I-D.ietf-mpls-bfd-directed]. For the case of MPLS dataplane, Segment Routing Architecture [I-D.ietf-spring-segment-routing] explains that "a segment is encoded as an MPLS label. An ordered list of segments is encoded as a stack of labels." Following on that this document defines Segment Routing with MPLS dataplane sub-TLV that MAY be used with the BFD Reverse Path TLV [I-D.ietf-mpls-bfd-directed]. The format of the sub-TLV is presented in Figure 1.

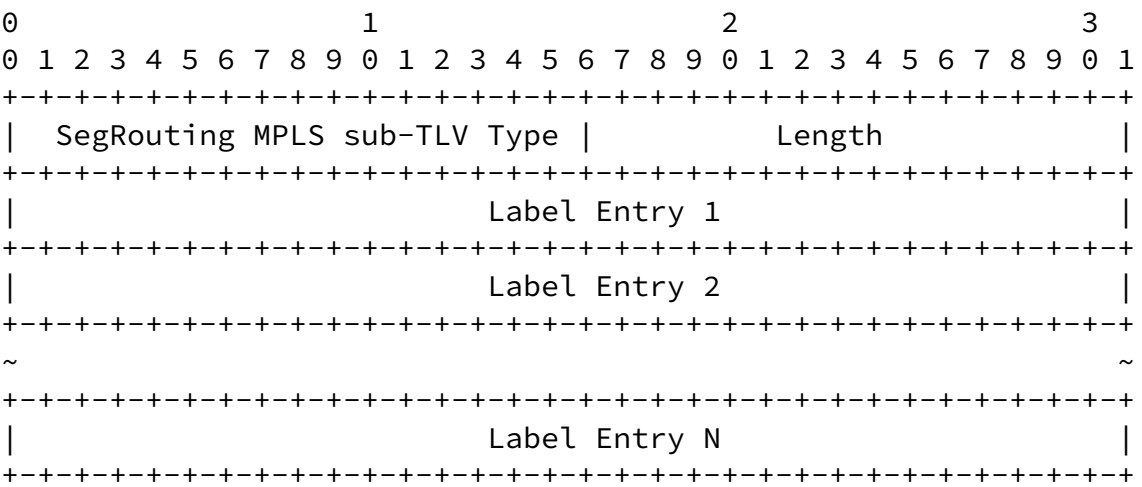

Figure 1: Segment Routing MPLS Tunnel sub-TLV

The Segment Routing Tunnel sub-TLV Type is two octets in length, and has a value of TBD (to be assigned by IANA as requested in Section 4).

The egress LSR MUST use the Value field as label stack for BFD control packets for the BFD session identified by the source IP address of the MPLS LSP Ping packet and the value in the BFD Discriminator TLV. Label Entries MUST be in network order.

Exactly one Segment Routing Tunnel sub-TLV MUST be included in the Reverse Path TLV. If more than one Segment Routing Tunnel sub-TLV is present in the Reverse Path TLV, then, in order to avoid ambiguity of which of TLVs to use, the egress BFD peer MUST send Echo Reply with the received Reverse Path TLVs and set the Return Code to "Too Many TLVs Detected" [I-D.ietf-mpls-bfd-directed]

The Segment Routing Tunnel sub-TLV MAY be used in Reply Path TLV defined in [RFC7110]

4. IANA Considerations

The IANA is requested to assign new sub-TLV type from "Multiprotocol Label Switching Architecture (MPLS) Label Switched Paths (LSPs) Ping Parameters - TLVs" registry, "Sub-TLVs for TLV Types 1, 16, and 21" sub-registry.

Value	Description	Reference
X (TBD)	Segment Routing MPLS Tunnel sub-TLV	This document

Table 1: New Segment Routing Tunnel sub-TLV

5. Security Considerations

Security considerations discussed in [RFC5880], [RFC5884], [RFC7726], and [RFC8029] apply to this document.

6. Acknowledgements

TBD

7. Normative References

[I-D.ietf-mpls-bfd-directed]
Mirsky, G., Tantsura, J., Varlashkin, I., and M. Chen, "Bidirectional Forwarding Detection (BFD) Directed Return Path", draft-ietf-mpls-bfd-directed-06 (work in progress), April 2017.

[I-D.ietf-mpls-spring-lsp-ping]
Kumar, N., Swallow, G., Pignataro, C., Akiya, N., Kini, S., Gredler, H., and M. Chen, "Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane", draft-ietf-mpls-spring-lsp-ping-02 (work in progress), December 2016.

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", draft-ietf-spring-segment-routing-11 (work in progress), February 2017.

[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>.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <http://www.rfc-editor.org/info/rfc5880>.

[RFC5881] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, DOI 10.17487/RFC5881, June 2010, <http://www.rfc-editor.org/info/rfc5881>.

[RFC5883] Katz, D. and D. Ward, "Bidirectional Forwarding Detection

(BFD) for Multihop Paths", RFC 5883, DOI 10.17487/RFC5883, June 2010, <http://www.rfc-editor.org/info/rfc5883>.

[RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, DOI 10.17487/RFC5884, June 2010, <http://www.rfc-editor.org/info/rfc5884>.

[RFC6428] Allan, D., Ed., Swallow, G., Ed., and J. Drake, Ed., "Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile", RFC 6428, DOI 10.17487/RFC6428, November 2011, <http://www.rfc-editor.org/info/rfc6428>.

[RFC7110] Chen, M., Cao, W., Ning, S., Jounay, F., and S. Delord, "Return Path Specified Label Switched Path (LSP) Ping", RFC 7110, DOI 10.17487/RFC7110, January 2014, <http://www.rfc-editor.org/info/rfc7110>.

[RFC7726] Govindan, V., Rajaraman, K., Mirsky, G., Akiya, N., and S. Aldrin, "Clarifying Procedures for Establishing BFD Sessions for MPLS Label Switched Paths (LSPs)", RFC 7726, DOI 10.17487/RFC7726, January 2016, <http://www.rfc-editor.org/info/rfc7726>.

[RFC8029] Kompella, K., Swallow, G., Pignataro, C., Ed., Kumar, N., Aldrin, S., and M. Chen, "Detecting Multiprotocol Label Switched (MPLS) Data-Plane Failures", RFC 8029, DOI 10.17487/RFC8029, March 2017, <http://www.rfc-editor.org/info/rfc8029>.

Author's Address

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com