

MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 18, 2017

G. Mirsky
Ericsson
J. Tantsura
Individual
I. Varlashkin
Google
M. Chen
Huawei
September 14, 2016

**Bidirectional Forwarding Detection (BFD) Directed Return Path
draft-ietf-mpls-bfd-directed-04**

Abstract

Bidirectional Forwarding Detection (BFD) is expected to be able to monitor wide variety of encapsulations of paths between systems. When a BFD session monitors an explicitly 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.

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 March 18, 2017.

Copyright Notice

Copyright (c) 2016 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
1.1.	Conventions used in this document	3
1.1.1.	Requirements Language	3
2.	Problem Statement	3
3.	Direct Reverse BFD Path	3
3.1.	Case of MPLS Data Plane	3
3.1.1.	BFD Reverse Path TLV	4
3.1.2.	Static and RSVP-TE sub-TLVs	5
3.1.3.	Segment Routing: MPLS Data Plane Case	5
3.2.	Bootstrapping BFD session with BFD Reverse Path over Segment Routed tunnel	5
3.3.	Return Codes	6
4.	Use Case Scenario	6
5.	IANA Considerations	7
5.1.	TLV	7
5.2.	Sub-TLV	7
5.3.	Return Codes	8
6.	Security Considerations	8
7.	Acknowledgements	8
8.	References	8
8.1.	Normative References	9
8.2.	Informative References	10
	Authors' Addresses	10

1. Introduction

[RFC 5880](#) [[RFC5880](#)], [RFC 5881](#) [[RFC5881](#)], and [RFC 5883](#) [[RFC5883](#)] established the BFD protocol for IP networks and [RFC 5884](#) [[RFC5884](#)] set rules of using BFD asynchronous mode over IP/MPLS LSPs. These standards implicitly assume that the egress BFD peer will use the shortest path route regardless of route being used to send BFD control packets towards it.

For the case where a LSP is explicitly routed it is likely that the shortest return path to the ingress BFD peer would not follow the same path as the LSP in the forward direction. The fact that BFD control packets are not guaranteed to follow the same links and nodes in both forward and reverse directions is a significant factor in producing false positive defect notifications, i.e. false alarms, if

used by the ingress BFD peer to deduce the state of the forward direction.

This document defines the BFD Reverse Path TLV as an extension to LSP Ping [[RFC4379](#)] and proposes that it is to be used to instruct the egress BFD peer to use explicit path for its BFD control packets associated with a particular BFD session. The TLV will be allocated from the TLV and sub-TLV registry defined by [RFC 4379](#) [[RFC4379](#)]. As a special case, forward and reverse directions of the BFD session can form a bi-directional co-routed associated channel.

1.1. Conventions used in this document

1.1.1. 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. Problem Statement

When BFD is used to monitor unidirectional explicitly routed path, e.g. MPLS-TE LSP, BFD control packets in forward direction would be in-band using the mechanism defined in [[RFC5884](#)] and [[RFC5586](#)]. But the reverse direction of the BFD session would follow the shortest path route and that might lead to the problem in detecting failures on a unidirectional explicit path as described below:

- o a failure detection by ingress node on the reverse path cannot be interpreted as bi-directional failure unambiguously and thus trigger, for example, protection switchover of the forward direction without possibility of being a false positive.

To address this scenario the egress BFD peer would be instructed to use a specific path for BFD control packets.

3. Direct Reverse BFD Path

3.1. Case of MPLS Data Plane

LSP ping, defined in [[RFC4379](#)], uses BFD Discriminator TLV [[RFC5884](#)] to bootstrap a BFD session over an MPLS LSP. This document defines a new TLV, BFD Reverse Path TLV, that MUST contain a single sub-TLV that can be used to carry information about the reverse path for the BFD session that is specified by value in BFD Discriminator TLV.

3.1.1. BFD Reverse Path TLV

The BFD Reverse Path TLV is an optional TLV within the LSP ping [RFC4379], [RFC6424]. However, if used, the BFD Discriminator TLV MUST be included in an Echo Request message as well. If the BFD Discriminator TLV is not present when the BFD Reverse Path TLV is included, then it MUST be treated as malformed Echo Request, as described in [RFC4379].

The BFD Reverse Path TLV carries information about the path onto which the egress BFD peer of the BFD session referenced by the BFD Discriminator TLV MUST transmit BFD control packets. The format of the BFD Reverse Path TLV is as presented in Figure 1.




Figure 1: BFD Reverse Path TLV

BFD Reverse Path TLV Type is 2 octets in length and has a value of TBD1 (to be assigned by IANA as requested in [Section 5](#)).

Length field is 2 octets long and defines the length in octets of the Reverse Path field.

Reverse Path field contains a sub-TLV. Any Target FEC sub-TLV (already defined, or to be defined in the future) for TLV Types 1, 16, and 21 of MPLS LSP Ping Parameters registry MAY be used in this field. Exactly one sub-TLV MUST be included in the Reverse Path TLV. If more than one 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" [Section 3.3](#).

If the egress LSR cannot find the path specified in the Reverse Path TLV it MUST send Echo Reply with the received Reverse Path TLV and set the Return Code to "Failed to establish the BFD session. The specified reverse path was not found" [Section 3.3](#). The egress BFD peer MAY establish the BFD session over IP network as defined in [RFC5884].

3.1.2. Static and RSVP-TE sub-TLVs

When an explicit path on an MPLS data plane is set either as Static or RSVP-TE LSP respective sub-TLVs defined in [RFC7110] MAY be used to identify the explicit reverse path for the BFD session.

3.1.3. Segment Routing: MPLS Data Plane Case

In addition to Static and RSVP-TE, Segment Routing with MPLS data plane can be used to set an explicit path. In this case a new sub-TLV is defined in this document as presented in Figure 2.



Figure 2: Segment Routing MPLS Tunnel sub-TLV

The Segment Routing Tunnel sub-TLV Type is two octets in length, and has a value of TBD2 (to be assigned by IANA as requested in [Section 5](#)).

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.

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

3.2. Bootstrapping BFD session with BFD Reverse Path 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 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. "

When LSP ping is used to bootstrap BFD session this document updates the statement and defines that LSP Ping MUST include the 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.

3.3. Return Codes

This document defines the following Return Codes for MPLS LSP Echo Reply:

- o "Too Many TLVs Detected", (TBD3). When more than one Reverse Path TLV found in the received Echo Request by the egress BFD peer, an Echo Reply with the return code set to "Too Many TLVs Detected" MUST be sent to the ingress BFD peer [Section 3.1.1](#).
- o "Failed to establish the BFD session. The specified reverse path was not found", (TBD4). When a specified reverse path is not available at the egress BFD peer, an Echo Reply with the return code set to "Failed to establish the BFD session. The specified reverse path was not found" MUST be sent back to the ingress BFD peer [Section 3.1.1](#).

4. Use Case Scenario

In the network presented in Figure 3 node A monitors two tunnels to node H: A-B-C-D-G-H and A-B-E-F-G-H. To bootstrap a BFD session to monitor the first tunnel, node A MUST include a BFD Discriminator TLV with Discriminator value (e.g. foobar-1) and MAY include a BFD Reverse Path TLV that references H-G-D-C-B-A tunnel. To bootstrap a BFD session to monitor the second tunnel, node A MUST include a BFD Discriminator TLV with a different Discriminator value (e.g. foobar-2) [\[RFC7726\]](#) and MAY include a BFD Reverse Path TLV that references H-G-F-E-B-A tunnel.




Figure 3: Use Case for BFD Reverse Path TLV

If an operator needs node H to monitor a path to node A, e.g. H-G-D-C-B-A tunnel, then by looking up list of known Reverse Paths it MAY find and use the existing BFD session.

5. IANA Considerations

5.1. TLV

The IANA is requested to assign a new value for BFD Reverse Path TLV from the "Multiprotocol Label Switching Architecture (MPLS) Label Switched Paths (LSPs) Ping Parameters - TLVs" registry, "TLVs and sub-TLVs" sub-registry.

Value	Description	Reference
X (TBD1)	BFD Reverse Path TLV	This document

Table 1: New BFD Reverse Type TLV

5.2. Sub-TLV

The IANA is requested to create new sub-registry for sub-TLV types of TLV TBD. All code points in the ranges 0 through 16383 and 32768 through 49161 in this registry shall be allocated according to the "IETF Review" procedure as specified in [RFC5226]. Code points in the ranges 16384 through 31743 and 49162 through 64511 in this registry shall be allocated according to the "First Come First Served" procedure as specified in [RFC5226]. Values in the range 31744 through 32767 and 64512 through 65534 are for Vendor or Private Use, and MUST NOT be allocated. This document defines the following new values of new sub-TLV type:

Value	Description	Reference
0	Reserved	This document
X (TBD2)	Segment Routing MPLS Tunnel sub-TLV	This document
2-31743	Unassigned	
31744-32767	Reserved for Vendor or Private Use	
32768-64511	Unassigned	
64512-65534	Reserved for Vendor or Private Use	
65535	Reserved	This document

Table 2: New Segment Routing Tunnel sub-TLV

5.3. Return Codes

The IANA is requested to assign a new Return Code value from the "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry, "Return Codes" sub-registry, as follows using a Standards Action value.

Value	Description	Reference
X (TBD3)	Too Many TLVs Detected.	This document
X (TBD4)	Failed to establish the BFD session.	This document
	The specified reverse path was not found.	

Table 3: New Return Code

6. Security Considerations

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

7. Acknowledgements

Authors greatly appreciate thorough review and the most helpful comments from Eric Gray and Carlos Pignataro.

8. References

8.1. Normative References

- [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-00](#) (work in progress), May 2016.
- [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>>.
- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", [RFC 4379](#), DOI 10.17487/RFC4379, February 2006, <<http://www.rfc-editor.org/info/rfc4379>>.
- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", [RFC 5586](#), DOI 10.17487/RFC5586, June 2009, <<http://www.rfc-editor.org/info/rfc5586>>.
- [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>>.
- [RFC6424] Bahadur, N., Kompella, K., and G. Swallow, "Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels", [RFC 6424](#), DOI 10.17487/RFC6424, November 2011, <<http://www.rfc-editor.org/info/rfc6424>>.

- [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>>.

8.2. Informative References

- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), DOI 10.17487/RFC5226, May 2008, <<http://www.rfc-editor.org/info/rfc5226>>.

Authors' Addresses

Greg Mirsky
Ericsson

Email: gregimirsky@gmail.com

Jeff Tantsura
Individual

Email: jefftant.ietf@gmail.com

Ilya Varlashkin
Google

Email: Ilya@nobulus.com

Mach(Guoyi) Chen
Huawei

Email: mach.chen@huawei.com

