

MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 8, 2018

G. Mirsky
ZTE
J. Tantsura
Individual
I. Varlashkin
Google
M. Chen
Huawei
December 5, 2017

**Bidirectional Forwarding Detection (BFD) Directed Return Path
draft-ietf-mpls-bfd-directed-08**

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 a 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 <https://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 June 8, 2018.

Copyright Notice

Copyright (c) 2017 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 (<https://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.	Control of the Reverse BFD Path	3
3.1.	BFD Reverse Path TLV	3
3.2.	Static and RSVP-TE sub-TLVs	5
3.3.	Return Codes	5
4.	Use Case Scenario	5
5.	IANA Considerations	6
5.1.	BFD Reverse Path TLV	6
5.2.	Return Code	6
6.	Security Considerations	6
7.	Normative References	6
Appendix A.	Acknowledgments	8
	Authors' Addresses	8

1. Introduction

[RFC5880], [RFC5881], and [RFC5883] established the BFD protocol for IP networks. [RFC5884] and [RFC7726] 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 [RFC8029] and proposes that it is to be used to instruct the egress BFD peer to use an 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 in [RFC8029]. 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 [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

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. Control of the Reverse BFD Path

LSP ping, defined in [[RFC8029](#)], 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. BFD Reverse Path TLV

The BFD Reverse Path TLV is an optional TLV within the LSP ping [[RFC8029](#)]. 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 [[RFC8029](#)].

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 non-multicast Target FEC Stack 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. Multicast Target FEC Stack sub-TLVs, i.e. p2mp and mp2mp, SHOULD NOT be included into Reverse Path field. If the egress LSR finds multicast Target Stack sub-TLV it MUST send echo reply with the received Reverse Path TLV, BFD Discriminator TLV and set the Return Code to "Inappropriate Target FEC Stack sub-TLV present" [Section 3.3](#). None, one or more sub-TLVs MAY be included in the BFD Reverse Path TLV. If none sub-TLVs found in the BFD Reverse Path TLV, the egress BFD peer MUST revert to using the local policy based decision as described in [Section 7 \[RFC5884\]](#), i.e., routed over IP network.

If the egress LSR cannot find the path specified in the Reverse Path TLV it MUST send Echo Reply with the received BFD Discriminator TLV, 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 [\[RFC5884\]](#).

3.2. 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.

3.3. Return Codes

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

- o "Inappropriate Target FEC Stack sub-TLV present", (TBD3). When multicast Target FEC Stack sub-TLV found in the received Echo Request by the egress BFD peer, an Echo Reply with the return code set to "Inappropriate Target FEC Stack sub-TLV present" MUST be sent to the ingress BFD peer [Section 3.1](#).
- 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](#).

4. Use Case Scenario

In the network presented in Figure 2 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 2: 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. BFD Reverse Path 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
(TBD1)	BFD Reverse Path TLV	This document

Table 1: New BFD Reverse Type TLV

5.2. Return Code

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
(TBD3)	Inappropriate Target FEC Stack sub-TLV present.	This document
(TBD4)	Failed to establish the BFD session. The specified reverse path was not found.	This document

Table 2: New Return Code

6. Security Considerations

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

7. Normative References

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

- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", [RFC 5586](#), DOI 10.17487/RFC5586, June 2009, <<https://www.rfc-editor.org/info/rfc5586>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", [RFC 5880](#), DOI 10.17487/RFC5880, June 2010, <<https://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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc5884>>.
- [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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc8029>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[Appendix A](#). Acknowledgments

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

Authors' Addresses

Greg Mirsky
ZTE

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

