

Workgroup: MPLS Working Group

Internet-Draft:

draft-ietf-mpls-bfd-directed-31

Published: 13 May 2024

Intended Status: Experimental

Expires: 14 November 2024

Authors: G. Mirsky    J. Tantsura    I. Varlashkin    M. Chen

Ericsson    NVIDIA    Google    Huawei

**Bidirectional Forwarding Detection (BFD) Directed Return Path for MPLS  
Label Switched Paths (LSPs)**

**Abstract**

Bidirectional Forwarding Detection (BFD) is expected to be able to monitor a 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 the egress BFD peer to use a specific path for the reverse direction of the BFD session. This document describes an extension to the MPLS Label Switched Path (LSP) echo request that allows a BFD system to request that the remote BFD peer transmits BFD control packets over the specified LSP.

**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 14 November 2024.

**Copyright Notice**

Copyright (c) 2024 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

## Table of Contents

- [1. Introduction](#)
  - [1.1. Conventions used in this document](#)
    - [1.1.1. Terminology](#)
    - [1.1.2. Requirements Language](#)
- [2. Problem Statement](#)
- [3. Control of the Reverse BFD Path](#)
  - [3.1. BFD Reverse Path TLV](#)
  - [3.2. Return Codes](#)
  - [3.3. Failure Characterization](#)
- [4. Use Case Scenario](#)
- [5. Operational Considerations](#)
- [6. IANA Considerations](#)
  - [6.1. BFD Reverse Path TLV](#)
  - [6.2. Return Code](#)
- [7. Implementation Status](#)
- [8. Security Considerations](#)
- [9. Normative References](#)
- [10. Informative References](#)
- [Appendix A. Acknowledgments](#)
- [Authors' Addresses](#)

## 1. Introduction

[[RFC5880](#)], [[RFC5881](#)], and [[RFC5883](#)] established the Bidirectional Forwarding Detection (BFD) protocol for IP networks. [[RFC5884](#)] and [[RFC7726](#)] set rules for using BFD Asynchronous mode over MPLS Label Switched Paths (LSPs), while not defining means to control the path an egress BFD system uses to send BFD control packets towards the ingress BFD system.

When BFD is used to detect defects of the traffic-engineered LSP, the path of the BFD control packets transmitted by the egress BFD system toward the ingress may be disjoint from the monitored 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 may be one of the factors contributing to producing false positive defect notifications, i.e., false alarms, at the ingress BFD peer. Ensuring that both directions of the BFD session use co-routed paths may, in some environments, improve the determinism of the failure detection and localization.

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 system 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.

The LSP ping extension, described in this document, was developed and implemented resulting from the operational experiment. The lessons learned from the operational experiment enabled the use between systems conforming to this specification. More implementations are encouraged to understand better the operational impact of the mechanism described in the document.

## **1.1. Conventions used in this document**

### **1.1.1. Terminology**

BFD: Bidirectional Forwarding Detection

FEC: Forwarding Equivalency Class

LSP: Label Switched Path

LSR: Label-Switching 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 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 an explicitly routed unidirectional path, e.g., MPLS-TE LSP, BFD control packets in the forward direction would be in-band using the mechanism defined in [[RFC5884](#)]. However, the reverse direction of the BFD session would follow the shortest path route, which could be completely or partially disjoint from the forward path. This creates the potential for the failure of a disjoint resource on the reverse path to trigger a BFD failure detection, even though the forward path is unaffected.

If the reverse path is congruent with the forward path, the potential for such false positives is greatly reduced. For this purpose, this specification provides a means for the egress BFD peer to be instructed to use a specific path for BFD control packets.

### 3. Control of the Reverse BFD Path

To bootstrap a BFD session over an MPLS LSP, LSP ping, defined in [RFC8029], MUST be used with BFD Discriminator TLV [RFC5884]. This document defines a new TLV, BFD Reverse Path TLV, that MAY contain none, one or more sub-TLVs that can be used to carry information about the reverse path for the BFD session that is specified by the 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 two octets in length and has a value of TBD1 (to be assigned by IANA as requested in Section 6).

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

Reverse Path field contains none, one, or more sub-TLVs. Only non-multicast Target FEC Stack sub-TLVs (already defined, or to be defined in the future) for TLV Types 1, 16, and 21 of MPLS LSP Ping Parameters registry are permitted to be used in this field. Any other sub-TLV MUST NOT be used. (This implies that Multicast Target FEC Stack sub-TLVs, i.e., p2mp and mp2mp, are not permitted in the Reverse Path field.) If the egress Label-Switching Router (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 3.2](#)). The BFD Reverse Path TLV includes none, one or more sub-TLVs. However, the number of sub-TLVs in the Reverse Path field MUST be limited. The default limit is 128 sub-TLV entries, but an implementation MAY be able to control that limit. An empty BFD Reverse Path TLV, i.e., no sub-TLVs present, is used as withdrawal of any previously set reverse path for the BFD session identified in the BFD Discriminator TLV. If no sub-TLVs are 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 of [[RFC5884](#)], i.e., routed over IP network.

If the egress peer 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 3.2](#)). If an implementation provides additional configuration options, these can control actions at the egress BFD peer, including when the path specified in the Reverse Path TLV cannot be found. For example, optionally, if the egress peer LSR cannot find the path specified in the Reverse Path TLV, it MAY establish the BFD session over an IP network, as defined in [[RFC5884](#)]. Note that the return code required by the MUST clause does not preclude the session from being established over a different path as discussed in the MAY clause.

The BFD Reverse Path TLV MAY be used in the bootstrapping of a BFD session process described in Section 6 of [[RFC5884](#)]. A system that supports this specification MUST support using the BFD Reverse Path TLV after the BFD session has been established. If a system that supports this specification receives an LSP Ping with the BFD Discriminator TLV and no BFD Reverse Path TLV even though the reverse path for the specified BFD session has been established according to the previously received BFD Reverse Path TLV, the egress BFD peer MUST transition to transmitting periodic BFD Control messages as defined in Section 7 of [[RFC5884](#)]. If a BFD system that received an LSP Ping with the BFD Reverse Path TLV does not support this specification, it will "result in the Return Code of 2 ("One or more of the TLVs was not understood") being sent in the echo response" (Section 3 of [[RFC8029](#)]).

### 3.2. Return Codes

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

- \*"Inappropriate Target FEC Stack sub-TLV present" (TBD3). When multicast Target FEC Stack sub-TLV found in the received Echo

Request, the egress BFD peer sends an Echo Reply with the return code set to "Inappropriate Target FEC Stack sub-TLV present" to the ingress BFD peer [Section 3.1](#).

\*"Failed to establish the BFD session. The specified reverse path was not found" (TBD4). When a specified reverse path is unavailable, the egress BFD peer sends an Echo Reply with the return code set to "Failed to establish the BFD session. The specified reverse path was not found" to the ingress BFD peer [Section 3.1](#).

### 3.3. Failure Characterization

A failure detected by a BFD session that uses the BFD Reverse Path TLV could be due to a change in the FEC used in the BFD Reverse Path TLV. The ingress BFD peer, upon detection of the network failure, MUST transmit the LSP Ping Echo request with Return Path TLV to verify whether the FEC is still valid. If the failure was caused by the change in the FEC used for the reverse direction of the BFD session, the ingress BFD peer MUST re-direct the reverse path of the BFD session using another FEC in BFD Reverse Path TLV, and notify an operator.

## 4. Use Case Scenario

In the network presented in [Figure 2](#), ingress LSR peer A monitors two tunnels to the egress LSR peer H: A-B-C-D-G-H and A-B-E-F-G-H. To bootstrap a BFD session to monitor the first tunnel, the ingress LSR peer A includes a BFD Discriminator TLV with Discriminator value (e.g., foobar-1). Peer A includes a BFD Reverse Path TLV referencing the H-G-D-C-B-A tunnel to control the path from the egress LSR. To bootstrap a BFD session to monitor the second tunnel, ingress LSR peer A, includes a BFD Discriminator TLV with a different Discriminator value (e.g., foobar-2) [[RFC7726](#)] and a BFD Reverse Path TLV that references the H-G-F-E-B-A tunnel.




Figure 2: Use Case for BFD Reverse Path TLV

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

## 5. Operational Considerations

When an explicit path is set either as Static or RSVP-TE LSP, corresponding sub-TLVs, defined in [RFC7110], MAY be used to identify the explicit reverse path for the BFD session. If a particular set of sub-TLVs composes the Return Path TLV [RFC7110] and does not increase the length of the Maximum Transmission Unit for the given LSP, that set can be safely used in the BFD Reverse Path TLV. If any of defined in [RFC7110] sub-TLVs used in BFD Reverse Path TLV, then the periodic verification of the control plane against the data plane, as recommended in Section 4 of [RFC5884], MUST use the Return Path TLV, as per [RFC7110], with that sub-TLV. By using the LSP Ping with Return Path TLV, an operator monitors whether at the egress BFD node the reverse LSP is mapped to the same FEC as the BFD session. Selection and control of the rate of LSP Ping with Return Path TLV follows the recommendation of [RFC5884]: "The rate of generation of these LSP Ping Echo request messages SHOULD be significantly less than the rate of generation of the BFD Control packets. An implementation MAY provide configuration options to control the rate of generation of the periodic LSP Ping Echo request messages."

Suppose an operator planned network maintenance activity that possibly affects FEC used in the BFD Reverse Path TLV. In that case, the operator can avoid the unnecessary disruption by using the LSP Ping with a new FEC in the BFD Reverse Path TLV. But in some scenarios, proactive measures cannot be taken. Because the frequency of LSP Ping messages will be lower than the defect detection time provided by the BFD session. As a result, a change in the reverse-path FEC will first be detected as the BFD session's failure. An operator will be notified as described in [Section 3.3](#).

## 6. IANA Considerations

### 6.1. BFD Reverse Path TLV

The IANA is requested to assign a new value for BFD Reverse Path TLV from the 16384-31739 range in the "TLVs" registry of "Multiprotocol Label Switching Architecture (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry.

| Type   | TLV Name             | Reference     | Sub-TLV Registry   |
|--------|----------------------|---------------|--|
| (TBD1) | BFD Reverse Path TLV | This document | Only non-multicast sub-TLV (already defined or to be defined in the future) at <a href="https://www.iana.org/assignments/mpls-lsp-ping-parameters/mpls-lsp-ping-parameters">[https://www.iana.org/assignments/mpls-lsp-ping-parameters/mpls-lsp-ping-parameters]</a> |

| Type | TLV Name | Reference | Sub-TLV Registry   |
|------|----------|-----------|--|
|      |          |           | <a href="#">parameters.xml#sub-tlv-1-16-21</a> are permitted to be used in this field. Any other sub-TLV MUST NOT be used. |

Table 1: New BFD Reverse Type TLV

## 6.2. Return Code

The IANA is requested to assign new Return Code values from the range 192-247 of the "Return Codes" registry of the "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters", as in [Table 2](#).

| 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

## 7. Implementation Status

Note to RFC Editor: This section MUST be removed before publication of the document.

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [\[RFC7942\]](#). The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [\[RFC7942\]](#), "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

- The organization responsible for the implementation: ZTE Corporation.



- The implementation's name ROSng empowers commonly used routers, e.g., ZXCTN 6000.
- A brief general description: A Return Path can be specified for a BFD session over RSVP tunnel or LSP. The same can be specified for a backup RSVP tunnel/LSP.

The implementation's level of maturity: production.

- Coverage: RSVP LSP (no support for Static LSP)
- Version compatibility: draft-ietf-mpls-bfd-directed-10.
- Licensing: proprietary.
- Implementation experience: simple once you support RFC 7110.
- Contact information: Qian Xin qian.xin2@zte.com.cn
- The date when information about this particular implementation was last updated: 12/16/2019

## 8. Security Considerations

Security considerations discussed in [[RFC5880](#)], [[RFC5884](#)], [[RFC7726](#)], [[RFC8029](#)], and [[RFC7110](#)] apply to this document.

BFD Reverse Path TLV may be exploited as an attack vector by inflating the number of included sub-TLVs. The number of sub-TLVs MUST be limited to mitigate that threat. The default limit for the number of sub-TLVs is set in [Section 3.1](#) to 128. An implementation MAY use a mechanism to control that limit.

## 9. 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>>.
- [[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>>.

## 10. Informative References

[RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <<https://www.rfc-editor.org/info/rfc7942>>.

## Appendix A. Acknowledgments

The authors greatly appreciate a thorough review and the most helpful comments from Eric Gray and Carlos Pignataro. The authors much appreciate the help of Qian Xin, who provided information about the implementation of this specification.

## Authors' Addresses

Greg Mirsky  
Ericsson

Email: [gregimirsky@gmail.com](mailto:gregimirsky@gmail.com)

Jeff Tantsura  
NVIDIA

Email: [jefftant.ietf@gmail.com](mailto:jefftant.ietf@gmail.com)

Ilya Varlashkin  
Google

Email: [imv@google.com](mailto:imv@google.com)

Mach(Guoyi) Chen  
Huawei

Email: [mach.chen@huawei.com](mailto:mach.chen@huawei.com)