

BFD for Multipoint Networks over Point-to-Multi-Point MPLS LSP
draft-mirsky-mpls-p2mp-bfd-06

Abstract

This document describes procedures for using Bidirectional Forwarding Detection (BFD) for multipoint networks to detect data plane failures in Multiprotocol Label Switching (MPLS) point-to-multipoint (p2mp) Label Switched Paths (LSPs). It also describes the applicability of LSP Ping, as in-band, and the control plane, as out-band, solutions to bootstrap a BFD session in this environment.

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 July 9, 2019.

Copyright Notice

Copyright (c) 2019 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
2.	Conventions used in this document	2
2.1.	Terminology	2
2.2.	Requirements Language	3
3.	Multipoint BFD Encapsulation	3
3.1.	IP Encapsulation of Multipoint BFD	3
3.2.	Non-IP Encapsulation of Multipoint BFD	4
4.	Bootstrapping Multipoint BFD	4
4.1.	LSP Ping	4
4.2.	Control Plane	5
5.	Security Considerations	5
6.	IANA Considerations	5
6.1.	Source MEP ID IP Address Type	5
7.	Acknowledgements	6
8.	References	6
8.1.	Normative References	6
8.2.	Informative References	7
	Author's Address	7

[1.](#) Introduction

[I-D.ietf-bfd-multipoint] defines a method of using Bidirectional Detection (BFD) [[RFC5880](#)] to monitor and detect unicast failures between the sender (head) and one or more receivers (tails) in multipoint or multicast networks. This document describes procedures for using such mode of BFD protocol to detect data plane failures in Multiprotocol Label Switching (MPLS) point-to-multipoint (p2mp) Label Switched Paths (LSPs). The document also describes the applicability of out-band solutions to bootstrap a BFD session in this environment.

[2.](#) Conventions used in this document

[2.1.](#) Terminology

MPLS: Multiprotocol Label Switching

LSP: Label Switched Path

BFD: Bidirectional Forwarding Detection

p2mp: Point-to-Multipoint

FEC: Forwarding Equivalence Class

G-ACh: Generic Associated Channel

ACH: Associated Channel Header

GAL: G-ACh Label

2.2. 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.

3. Multipoint BFD Encapsulation

[I-D.ietf-bfd-multipoint] defines how the tail of multipoint BFD session demultiplexes received BFD control packet when Your Discriminator is not set, i.e., equals zero. Because [[I-D.ietf-bfd-multipoint](#)] uses BFD in Demand mode from the very start of the p2mp BFD session, the head of BFD multipoint session transmits all BFD control packets with Your Discriminator set to zero. As a result, a tail cannot demultiplex BFD sessions using Your Discriminator, as defined in [[RFC5880](#)]. [[I-D.ietf-bfd-multipoint](#)] requires that to demultiplex BFD sessions the tail uses the source IP address, My Discriminator and the identity of the multipoint tree from which the Multipoint BFD Control packet was received. The identification of the multipoint tree MAY be provided by the p2mp MPLS LSP label in case of inclusive p-tree or upstream assigned label in case of aggregate p-tree. If the BFD control packet is encapsulated in IP/UDP, then the source IP address MUST be used to demultiplex the received BFD control packet as described in [Section 3.1](#). The non-IP encapsulation case is described in [Section 3.2](#).

3.1. IP Encapsulation of Multipoint BFD

[I-D.ietf-bfd-multipoint] defines IP/UDP encapsulation for multipoint BFD over p2mp MPLS LSP:

UDP destination port MUST be set to 3784;

destination IP address MUST be from the 127/8 range for IPv4 and from the 0:0:0:0:FFFF:7F00/104 range for IPv6;

This specification further clarifies that:

if multiple alternative paths for the given p2mp LSP Forwarding Equivalence Class (FEC) exist, the MultipointHead SHOULD use Entropy Label [RFC6790] used for LSP Ping [RFC8029] to exercise that particular alternative path;

or the MultipointHead MAY use the IP address discovered by LSP Ping traceroute [RFC8029] as the destination IP address to possibly exercise that particular alternate path.

3.2. Non-IP Encapsulation of Multipoint BFD

In some environments, the overhead of extra IP/UDP encapsulations may be considered as overburden and make using more compact G-ACh encapsulation attractive. Non-IP encapsulation for multipoint BFD over p2mp MPLS LSP MUST use Generic Associated Channel (G-ACh) Label (GAL) [RFC5586] at the bottom of the label stack followed by Associated Channel Header (ACH). Channel Type field in ACH MUST be set to MPLS-TP CV value (0x0023) [RFC6428]. To provide the identity of the MultipointHead for the particular multipoint BFD session this document defines new Source MEP ID IP Address type (TBA1) in [Section 6.1](#). If the Length value is 4, then the Value field contains an IPv4 address. If the Length value is 16, then the Value field contains an IPv6 address. Any other value of the Length field MUST be considered as an error, and the BFD control packet MUST be discarded.

4. Bootstrapping Multipoint BFD

4.1. LSP Ping

LSP Ping is the part of on-demand OAM toolset to detect and localize defects in the data plane, and verify the control plane against the data plane by ensuring that the LSP is mapped to the same FEC, at the egress, as the ingress.

LSP Ping, as defined in [RFC6425], MAY be used to bootstrap MultipointTail. If the LSP Ping used, it MUST include the Target FEC TLV and the BFD Discriminator TLV defined in [RFC5884]. The Target FEC TLV MUST use sub-TLVs defined in [Section 3.1](#) [RFC6425]. It is RECOMMENDED setting the value of Reply Mode field to "Do not reply" [RFC8029] for the LSP Ping to bootstrap MultipointTail of the p2mp BFD session. A MultipointTail that receives the LSP Ping that includes the BFD Discriminator TLV:

- o MUST validate the LSP Ping;
- o MUST associate the received BFD Discriminator value with the p2mp LSP;

- o MUST create p2mp BFD session and set `bfd.SessionType = MultipointTail` as described in [\[I-D.ietf-bfd-multipoint\]](#);
- o MUST use the source IP address of LSP Ping, the value of BFD Discriminator from the BFD Discriminator TLV, and the identity of the p2mp LSP to properly demultiplex BFD sessions.

Besides bootstrapping a BFD session over a p2mp LSP, LSP Ping SHOULD be used to verify the control plane against the data plane periodically by checking that the p2mp LSP is mapped to the same FEC at the MultipointHead and all active MultipointTails. 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 because LSP Ping requires more processing to validate the consistency between the data plane and the control plane. An implementation MAY provide configuration options to control the rate of generation of the periodic LSP Ping Echo request messages.

4.2. Control Plane

BGP-BFD Attribute [\[I-D.ietf-bess-mvpn-fast-failover\]](#) MAY be used to bootstrap multipoint BFD session on a tail.

5. Security Considerations

This document does not introduce new security aspects but inherits all security considerations from [\[RFC5880\]](#), [\[RFC5884\]](#), [\[RFC7726\]](#), [\[I-D.ietf-bfd-multipoint\]](#), [\[RFC8029\]](#), and [\[RFC6425\]](#).

Also, BFD for p2mp MPLS LSP MUST follow the requirements listed in [section 4.1 \[RFC4687\]](#) to avoid congestion in the control plane or the data plane caused by the rate of generating BFD control packets. An operator SHOULD consider the amount of extra traffic generated by p2mp BFD when selecting the interval at which the MultipointHead will transmit BFD control packets. Also, the operator MAY consider the size of the packet the MultipointHead transmits periodically as using IP/UDP encapsulation adds up to 28 octets, which is more than 50% of BFD control packet length, comparing to G-ACh encapsulation.

6. IANA Considerations

6.1. Source MEP ID IP Address Type

IANA is required to allocate value (TBD) for the Source MEP ID IP Address type from the "CC/CV MEP-ID TLV" registry which is under the "Pseudowire Associated Channel Types" registry.

Value	Description	Reference
TBA1	IP Address	This document

Table 1: Source MEP ID IP Address TLV Type

7. Acknowledgements

The author sincerely appreciates the comments received from Andrew Malis.

8. References

8.1. Normative References

- [I-D.ietf-bess-mvpn-fast-failover]
Morin, T., Kebler, R., and G. Mirsky, "Multicast VPN fast upstream failover", [draft-ietf-bess-mvpn-fast-failover-04](#) (work in progress), November 2018.
- [I-D.ietf-bfd-multipoint]
Katz, D., Ward, D., Networks, J., and G. Mirsky, "BFD for Multipoint Networks", [draft-ietf-bfd-multipoint-19](#) (work in progress), December 2018.
- [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>>.
- [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>>.

- [RFC6425] Saxena, S., Ed., Swallow, G., Ali, Z., Farrel, A., Yasukawa, S., and T. Nadeau, "Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping", [RFC 6425](#), DOI 10.17487/RFC6425, November 2011, <<https://www.rfc-editor.org/info/rfc6425>>.
- [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, <<https://www.rfc-editor.org/info/rfc6428>>.
- [RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", [RFC 6790](#), DOI 10.17487/RFC6790, November 2012, <<https://www.rfc-editor.org/info/rfc6790>>.
- [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>>.

8.2. Informative References

- [RFC4687] Yasukawa, S., Farrel, A., King, D., and T. Nadeau, "Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks", [RFC 4687](#), DOI 10.17487/RFC4687, September 2006, <<https://www.rfc-editor.org/info/rfc4687>>.

Author's Address

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com

