

Workgroup: BFD Working Group

Internet-Draft:

draft-mirsky-bfd-mpls-demand-10

Published: 1 October 2021

Intended Status: Standards Track

Expires: 4 April 2022

Authors: G. Mirsky

Ericsson

BFD in Demand Mode over Point-to-Point MPLS LSP

Abstract

This document describes procedures for using Bidirectional Forwarding Detection (BFD) in Demand mode to detect data plane failures in Multiprotocol Label Switching (MPLS) point-to-point Label Switched Paths.

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 4 April 2022.

Copyright Notice

Copyright (c) 2021 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. Conventions used in this document](#)
 - [2.1. Terminology](#)
 - [2.2. Requirements Language](#)
- [3. Use of the BFD Demand Mode](#)
 - [3.1. The Applicability of BFD for Multipoint Networks](#)
- [4. IANA Considerations](#)
- [5. Security Considerations](#)
- [6. Normative References](#)
- [7. Informative References](#)
- [Appendix A. Acknowledgements](#)
- [Author's Address](#)

1. Introduction

[[RFC5884](#)] defined use of the Asynchronous method of Bidirectional Detection (BFD) [[RFC5880](#)] to monitor and detect failures in the data path of a Multiprotocol Label Switching (MPLS) Label Switched Path (LSP). Use of the Demand mode, also specified in [[RFC5880](#)], has not been defined so far. This document describes procedures for using the Demand mode of BFD protocol to detect data plane failures in MPLS point-to-point (p2p) LSPs.

2. Conventions used in this document

2.1. Terminology

MPLS: Multiprotocol Label Switching

LSP: Label Switched Path

LER: Label switching Edge Router

BFD: Bidirectional Forwarding Detection

p2p: Point-to-Point

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. Use of the BFD Demand Mode

[[RFC5880](#)] defines that the Demand mode MAY be:

- *asymmetric, i.e. used in one direction of a BFD session;
- *switched to and from without bringing BFD session to Down state through using a Poll Sequence.

For the case of BFD over MPLS LSP, ingress Label switching Edge Router (LER) usually acts as Active BFD peer and egress LER acts as Passive BFD peer. The Active peer bootstraps the BFD session by using LSP ping. If the BFD session is configured to use the Demand mode, once the BFD session is in Up state the ingress LER MUST switch to the Demand mode as defined in Section 6.6 [[RFC5880](#)]. The egress LER also follows procedures defined in Section 6.6 [[RFC5880](#)] and ceases further transmission of periodic BFD control packets to the ingress LER.

In this state BFD peers MAY remain as long as the egress LER is in Up state. The ingress LER SHOULD periodically check continuity of a bidirectional path between the ingress and egress LERs by using the Poll Sequence, as described in Section 6.6 [[RFC5880](#)]. An implementation that supports using the Poll Sequence as the mechanism for bidirectional path continuity check MUST be able to control the interval between consecutive Poll Sequences. The RECOMMENDED default value is 1 second.

If the Detection timer at the egress LER expires it MUST send BFD Control packet to the ingress LER with the Poll (P) bit set, Status (Sta) field set to the Down value, and the Diagnostic (Diag) field set to Control Detection Time Expired value. The egress LER periodically transmits these Control packets to the ingress LER until either it receives the valid for this BFD session control packet with the Final (F) bit set from the ingress LER or the defect condition clears and the BFD session state reaches Up state at the egress LER. An implementation that supports this specification MUST provide control of the interval between consecutive Poll messages signaling the expiration of the Detection timer. The RECOMMENDED default value of the interval is 1 second.

The ingress LER transmits BFD Control packets over the MPLS LSP with the Demand (D) flag set at negotiated interval per [[RFC5880](#)], the greater of `bfd.DesiredMinTxInterval` and `bfd.RemoteMinRxInterval`, until it receives the valid BFD packet from the egress LER with the Poll (P) bit and the Diagnostic (Diag) field value Control Detection Time Expired. Reception of such BFD control packet by the ingress LER indicates that the monitored LSP has a failure and sending BFD control packet with the Final flag set to acknowledge failure

indication is likely to fail. Instead, the ingress LER transmits the BFD Control packet to the egress LER over the IP network with:

- *destination IP address MUST be set to the destination IP address of the LSP Ping Echo request message [[RFC8029](#)];

- *destination UDP port set to 4784 [[RFC5883](#)];

- *Final (F) flag in BFD control packet MUST be set;

- *Demand (D) flag in BFD control packet MUST be cleared.

The ingress LER changes the state of the BFD session to Down and changes rate of BFD Control packets transmission to one packet per second. The ingress LER in Down mode changes to Asynchronous mode until the BFD session comes to Up state once again. Then the ingress LER switches to the Demand mode.

3.1. The Applicability of BFD for Multipoint Networks

[[RFC8562](#)] defines the use of BFD in multipoint networks. This specification analyzes the case of p2p LSP. In that scenario, the ingress of the LSP acts as the MultipointHead, and the egress - as MultipointTail. The BFD state machines for MultipointHead, MultipointClient, and MultipointTail don't use the three-way handshakes for session establishment and teardown. As a result, the Init state is absent, and the session transitions to the Up state once the BFD session is administratively enabled. Hence, a BFD session over a p2p LSP, using principles of [[RFC8562](#)] or [[RFC8563](#)], can be established faster if the MultipointTail has been provisioned with the value of My Discriminator used by the MultipointHead for that BFD session. That value can be provided to the MultipointTail using different mechanisms, e.g., an extension to IGP. Description of mechanism to provide the value of My Discriminator used by the MultipointHead for the particular BFD session is outside the scope of this specification.

Unsolicited notification of the detected failure by the MultipointTail to the MultipointClient performs as described above for the case when the ingress BFD system switches the remote peer into the Demand mode.

4. IANA Considerations

TBD

5. Security Considerations

This document does not introduce new security aspects but inherits all security considerations from [RFC5880], [RFC5884], [RFC7726], [RFC8029], and [RFC6425].

6. 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>>.
- [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>>.
- [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>>.
- [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>>.

[RFC8562] Katz, D., Ward, D., Pallagatti, S., Ed., and G. Mirsky, Ed., "Bidirectional Forwarding Detection (BFD) for Multipoint Networks", RFC 8562, DOI 10.17487/RFC8562, April 2019, <<https://www.rfc-editor.org/info/rfc8562>>.

7. Informative References

[RFC8563] Katz, D., Ward, D., Pallagatti, S., Ed., and G. Mirsky, Ed., "Bidirectional Forwarding Detection (BFD) Multipoint Active Tails", RFC 8563, DOI 10.17487/RFC8563, April 2019, <<https://www.rfc-editor.org/info/rfc8563>>.

Appendix A. Acknowledgements

TBD

Author's Address

Greg Mirsky
Ericsson

Email: gregimirsky@gmail.com