

Internet Engineering Task Force (IETF)
Request for Comments: 7726
Updates: 5884
Category: Standards Track
ISSN: 2070-1721

V. Govindan
K. Rajaraman
Cisco Systems
G. Mirsky
Ericsson
N. Akiya
Big Switch Networks
S. Aldrin
Google
January 2016

Clarifying Procedures for Establishing BFD Sessions for
MPLS Label Switched Paths (LSPs)

Abstract

This document clarifies the procedures for establishing, maintaining, and removing multiple, concurrent BFD (Bidirectional Forwarding Detection) sessions for a given <MPLS LSP, FEC> as described in RFC 5884.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7726>.

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. Background 2
 - 1.1. Requirements Language 3
- 2. Theory of Operation 3
 - 2.1. Procedures for Establishment of Multiple BFD Sessions . . 3
 - 2.2. Procedures for Maintenance of Multiple BFD Sessions . . . 4
 - 2.3. Procedures for Removing BFD Sessions at the Egress LSR . 4
 - 2.4. Changing Discriminators for a BFD Session 5
- 3. Backwards Compatibility 5
- 4. Security Considerations 5
- 5. References 6
 - 5.1. Normative References 6
 - 5.2. Informative References 6
- Acknowledgements 6
- Authors' Addresses 7

1. Background

[RFC5884] defines the procedures to bootstrap and maintain BFD sessions for an <MPLS LSP, FEC> using a Label Switched Path (LSP) ping. While Section 4 of [RFC5884] specifies that multiple BFD sessions can be established for an <MPLS LSP, FEC> tuple, the procedures to bootstrap and maintain multiple BFD sessions concurrently over an <MPLS LSP, FEC> are not clearly specified. Additionally, the procedures of removing BFD sessions bootstrapped on the egress Label Switching Router (LSR) are unclear. This document provides those clarifications without deviating from the principles outlined in [RFC5884].

The ability for an ingress LSR to establish multiple BFD sessions for an <MPLS LSP, FEC> tuple is useful in scenarios such as LSPs based on Segment Routing [SEG-ROUTING] or LSPs having Equal-Cost Multipath (ECMP). The process used by the ingress LSR to determine the number of BFD session(s) to be bootstrapped for an <MPLS LSP, FEC> tuple and the mechanism used to construct those session(s) are outside the scope of this document.

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. Theory of Operation

2.1. Procedures for Establishment of Multiple BFD Sessions

Section 4 of [RFC5884] specifies the procedure for bootstrapping BFD sessions using LSP ping. It further states that "a BFD session SHOULD be established for each alternate path that is discovered." This requirement has been the source of some ambiguity as the procedures of establishing concurrent, multiple sessions have not been explicitly specified. This ambiguity can also be attributed in part to the text in Section 7 of [RFC5884] forbidding either end to change local discriminator values in BFD control packets after the session reaches the UP state. The following procedures are described to clarify the ambiguity based on the interpretation of the authors' reading of the referenced sections:

At the ingress LSR:

MPLS LSP ping can be used to bootstrap multiple BFD sessions for a given <MPLS LSP, FEC>. Each LSP ping MUST carry a different discriminator value in the BFD discriminator TLV [RFC5884].

The egress LSR needs to perform the following:

If the validation of the Forwarding Equivalence Class (FEC) in the MPLS Echo request message succeeds, check the discriminator specified in the BFD discriminator TLV of the MPLS Echo request. If there is no local session that corresponds to the (remote) discriminator received in the MPLS Echo request, a new session is bootstrapped and a local discriminator is allocated. The validation of a FEC is a necessary condition to be satisfied to create a new BFD session at the egress LSR. However, the policy or procedure, if any, to be applied by the egress LSR before

allowing a new BFD session to be created is outside the scope of this document. Such policies or procedures could consider availability of system resources before allowing a session to be created. When the egress LSR disallows the creation of a BFD session due to policy, it MUST drop the MPLS Echo request message.

Ensure the uniqueness of the <MPLS LSP, FEC, Remote Discriminator> tuple.

Except for the clarification mentioned above, the remaining procedures of BFD session establishment are as specified in Sections 4-6 of [RFC5884].

2.2. Procedures for Maintenance of Multiple BFD Sessions

Both the ingress LSR and egress LSR use the Your Discriminator of the received BFD packet to demultiplex BFD sessions.

2.3. Procedures for Removing BFD Sessions at the Egress LSR

[RFC5884] does not specify an explicit procedure for deleting BFD sessions. The procedure for removing a BFD session established by an out-of-band discriminator exchange using the MPLS LSP ping can improve resource management (e.g., memory), especially in scenarios involving thousands or more of such sessions. A few observations are made here:

The BFD session MAY be removed in the egress LSR if the BFD session transitions from UP to DOWN. This can either be done immediately after the BFD session transitions from UP to DOWN or after the expiry of a configurable timer started after the BFD session state transitions from UP to DOWN at the egress LSR to reduce flapping by adding hysteresis.

The BFD session on the egress LSR MAY be removed by the ingress LSR by using the BFD diagnostic code AdminDown(7) as specified in [RFC5880]. When the ingress LSR wants to remove a session without triggering any state change at the egress, it MAY transmit BFD packets indicating the State as Down(1), diagnostic code AdminDown(7) detectMultiplier number of times. Upon receiving such a packet, the egress LSR MAY remove the BFD session, without triggering a change of state.

The procedures to be followed at the egress LSR when BFD session(s) remain in the DOWN state for a significant amount of time is a local matter. Such procedures are outside the scope of this document.

All BFD sessions established with the FEC MUST be removed automatically if the FEC is removed.

The egress MUST use the discriminators exchanged when the session was brought UP to indicate any session state change to the ingress. The egress SHOULD reset this to zero after transmitting `bfd.detectMult` number of packets if the BFD session transitions to DOWN state.

2.4. Changing Discriminators for a BFD Session

The discriminators of a BFD session established over an MPLS LSP cannot be changed when it is in UP state. The BFD session could be removed after a graceful transition to AdminDown state using the BFD diagnostic code AdminDown. A new session could be established with a different discriminator. The initiation of the transition from the UP to DOWN state can be done by either the ingress LSR or the egress LSR.

3. Backwards Compatibility

The procedures clarified by this document are fully backward compatible with an existing implementation of [RFC5884]. While the capability to bootstrap and maintain multiple BFD sessions may not be present in current implementations, the procedures outlined by this document can be implemented as a software upgrade without affecting existing sessions. In particular, the egress LSR needs to support multiple BFD sessions per <MPLS LSP, FEC> before the ingress LSR is upgraded.

4. Security Considerations

This document clarifies the mechanism to bootstrap multiple BFD sessions per <MPLS LSP, FEC>. BFD sessions, naturally, use system and network resources. More BFD sessions means more resources will be used. It is highly important to ensure that only a minimum number of BFD sessions are provisioned per FEC and that bootstrapped BFD sessions are properly deleted when they are no longer required. Additionally, security measures described in [RFC4379] and [RFC5884] are to be followed.

5. References

5.1. 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, <<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>>.
- [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>>.
- [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>>.

5.2. Informative References

- [SEG-ROUTING] Filsfils, C., Ed., Previdi, S., Ed., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", Work in Progress, draft-ietf-spring-segment-routing-07, December 2015.

Acknowledgements

The authors would like to thank Marc Binderberger for performing thorough reviews and providing valuable suggestions.

The authors would like to thank Mudigonda Mallik, Rajaguru Veluchamy, and Carlos Pignataro of Cisco Systems for their review comments.

The authors would like to thank Alvaro Retana and Scott Bradner for their review comments.

Authors' Addresses

Vengada Prasad Govindan
Cisco Systems

Email: venggovi@cisco.com

Kalyani Rajaraman
Cisco Systems

Email: kalyanir@cisco.com

Gregory Mirsky
Ericsson

Email: gregory.mirsky@ericsson.com

Nobo Akiya
Big Switch Networks

Email: nobo.akiya.dev@gmail.com

Sam Aldrin
Google

Email: aldrin.ietf@gmail.com