

Internet Engineering Task Force
Internet-Draft
Updates: 5884 (if approved)
Intended status: Standards Track
Expires: April 6, 2015

V. Govindan
Cisco Systems
K. Rajaraman
G. Mirsky
Ericsson
N. Akiya
Cisco Systems
S. Aldrin
Huawei Technologies
October 3, 2014

Clarifications to RFC 5884
draft-grmas-bfd-rfc5884-clarifications-00

Abstract

This document clarifies the procedures for establishing, maintaining and removing multiple, concurrent BFD sessions for a given <MPLS LSP, FEC> described in RFC5884.

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 <http://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 April 6, 2015.

Copyright Notice

Copyright (c) 2014 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	2
2. Theory of Operation	3
2.1. Procedures for establishment of multiple BFD sessions . . .	3
2.2. Procedures for maintenance of multiple BFD sessions . . .	3
2.3. Procedures for removing BFD sessions at the egress LSR . .	4
2.4. Changing discriminators for a BFD session	4
3. Backwards Compatibility	4
4. Encapsulation	5
5. Security Considerations	5
6. IANA Considerations	5
7. Acknowledgements	5
8. Normative References	5
Authors' Addresses	5

1. Background

[RFC5884] defines the procedures to bootstrap and maintain BFD sessions for a <MPLS FEC, LSP> using LSP ping. While Section 4 of [RFC5884] specifies that multiple BFD sessions can be established for a <MPLS FEC, LSP> tuple, the procedures to bootstrap and maintain multiple BFD sessions concurrently over a <MPLS FEC, LSP> are not clearly specified. Additionally, the procedures of removing BFD sessions bootstrapped on the egress 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 a <MPLS FEC, LSP> tuple is useful in scenarios such as Segment Routing based LSPs 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 a <MPLS FEC, LSP> tuple and the mechanism of constructing 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 6 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's reading of the referenced sections:

At the ingress LSR:

MPLS LSP ping can be used to bootstrap multiple BFD sessions for a given <MPLS FEC, LSP>. Each LSP ping MUST carry a different discriminator value in the BFD discriminator TLV [RFC4379].

The egress LSR needs to perform the following:

If the validation of the 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 discriminator (remote) received in the MPLS Echo request, a new session is bootstrapped and a local discriminator is allocated. Since the BFD local discriminator of either ends cannot change as long as the session is in the UP state, a new discriminator received in the LSP ping unambiguously conveys the intent of the LSR ingress to bootstrap a new BFD session for the FEC specified in the LSP ping.

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

The remaining procedures of session establishment are as specified in [RFC5884].

2.2. Procedures for maintenance of multiple BFD sessions

Both the ingress LSR and egress LSR use the YourDiscriminator 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 (like memory etc.) especially in scenarios involving thousands or more of such sessions. A few options are possible here:

The BFD session MAY be removed in the egress LSR if the BFD session transitions from UP to DOWN. This can be done after the expiry of a configurable timer started after the BFD session state transitions from UP to DOWN at the egress LSR.

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

Ed Note: The procedures to be followed at the egress LSR when the BFD session never transitions to UP from DOWN state are yet to be clarified

Regardless of the option chosen to proceed, all BFD sessions established with the FEC MUST be removed automatically if the FEC is removed.

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 either by 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 FEC, LSP> before the ingress LSR is upgraded.

4. Encapsulation

The encapsulation of BFD packets are the same as specified by [RFC5884].

5. Security Considerations

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

6. IANA Considerations

This document does not make any requests to IANA.

7. Acknowledgements

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

8. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.
- [RFC5884] Aggarwal, R., Kompella, K., Nadeau, T., and G. Swallow, "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.

Authors' Addresses

Vengada Prasad Govindan
Cisco Systems

Email: venggovi@cisco.com

Kalyani Rajaraman
Ericsson

Email: kalyani.rajaraman@ericsson.com

Gregory Mirsky
Ericsson

Email: gregory.mirsky@ericsson.com

Nobo Akiya
Cisco Systems

Email: nobo@cisco.com

Sam Aldrin
Huawei Technologies

Email: aldrin.ietf@gmail.com