

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 5, 2015

V. Govindan
N. Akiya
Cisco Systems
July 4, 2014

**Label Switched Path (LSP) Ping
Extended Bidirectional Forwarding Detection (BFD) Discriminator TLV
draft-vgovindan-mpls-extended-bfd-disc-tlv-00**

Abstract

This document defines an extended Bidirectional Forwarding Detection (BFD) discriminator TLV for the Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Ping mechanism, to allow bootstrapping of multiple BFD sessions for a given FEC.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 January 5, 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
2.	Overview	3
3.	Procedures for BFD session establishment and removal using the Extended BFD TLV	3
3.1.	Procedures for establishing BFD sessions	3
3.2.	Procedures for removing BFD sessions	3
4.	Extended BFD Discriminator TLV	4
5.	Mutually Exclusive: BFD TLVs	5
6.	Backwards Compatibility	5
7.	Encapsulation	5
8.	Security Considerations	5
9.	IANA Considerations	6
9.1.	Extended BFD Discriminator TLV	6
10.	Acknowledgements	6
11.	Contributing Authors	6
12.	Normative References	6
Appendix A.	Alternate format for the BFD Extended TLV	6
	Authors' Addresses	8

[1.](#) Background

Bidirectional Forwarding Detection (BFD) [[RFC5880](#)] for Multiprotocol Label Switching (MPLS) Label Switched Paths (LSPs), [[RFC5884](#)], describes a mechanism to use BFD to monitor the connectivity in-band on the LSPs. The BFD session on the LSP egress is bootstrapped using the LSP Ping mechanism, defined in [[RFC4379](#)], carrying the BFD Discriminator TLV that describes the BFD discriminator of the BFD session on the LSP ingress.

The BFD Discriminator TLV and defined procedures around this TLV only allow one BFD session to be bootstrapped per <MPLS Forwarding Equivalent Class (FEC), LSP>. There are scenarios where an LSP ingress may desire to run multiple BFD sessions to monitor the connectivity on an LSP. To achieve the bootstrapping of multiple BFD sessions per FEC, a new TLV and procedures are required. Two scenarios where this is useful are described below:

- o Entropy labels help achieve load balancing of traffic belonging to the same <MPLS FEC, LSP>. It may be beneficial to track the

individual paths of the multi-path network using separate BFD sessions for each non-congruent path.

- o It may be useful to establish multiple BFD sessions for the same <MPLS FEC, LSP> to achieve BFD session redundancy, i.e. protection against false positives due to equipment or soft failures inside boxes.

2. Overview

An LSR ingress wanting to bootstrap one or more BFD sessions on an LSP is to include the Extended BFD Discriminator TLV, described in [Section 4](#), in the MPLS echo request message for the FEC. The Extended BFD Discriminator TLV is capable of carrying multiple BFD discriminators, and each BFD discriminator is accompanied with an instance identifier. The LSR egress, upon reception of this MPLS echo request, is to create requested number of BFD sessions for the specified FEC. Each BFD session object created on the LSR ingress and the LSR egress MUST be annotated with corresponding instance identifier. BFD session procedures are to follow those described in [\[RFC5884\]](#).

3. Procedures for BFD session establishment and removal using the Extended BFD TLV

3.1. Procedures for establishing BFD sessions

There are at least two options possible here:

1. BFD session establishment MUST follow the procedure specified in [\[RFC5884\]](#).
2. The base procedure for BFD session establishment MUST be the same as that of [\[RFC5884\]](#). This procedure can be enhanced by specifying additional Operation type field and Operation status field in the proposed Extended BFD Discriminator TLV. See [Appendix A](#) for a description of Operation types and Operation status codes.

3.2. Procedures for removing BFD sessions

[\[RFC5884\]](#) does not specify an explicit procedure for deleting BFD sessions. A few options are possible here:

1. Specify an explicit delete procedure for the BFD session using Operation types field and Operation status field through the Extended BFD TLV. See [Appendix A](#) for a description of Operation types and Operation status codes.

Instance Identifier - An instance identifier of the BFD session. The instance identifier is a value allocated by the LSP ingress

for corresponding BFD Discriminator, and MUST be unique within the FEC on the LSP ingress node. The instance identifier MUST NOT change for the lifetime of the BFD session.

BFD Discriminator - The BFD discriminator allocated for this BFD session by the LSP ingress.

See [Appendix A](#) for a discussion on an alternate format for the TLV.

5. Mutually Exclusive: BFD TLVs

The BFD Discriminator TLV and the Extended BFD Discriminator TLV are mutually exclusive. An MPLS echo request/reply message MUST NOT include both the BFD Discriminator TLV and the Extended BFD Discriminator TLV. Reception of an MPLS echo request with both the BFD Discriminator TLV and the Extended BFD Discriminator TLV is to result in the Return Code being set to Malformed echo request received (1).

6. Backwards Compatibility

If an LSP ingress wishes to bootstrap multiple BFD sessions with the Extended BFD Discriminator TLV when an LSP already has a BFD session bootstrapped with the BFD Discriminator TLV, following procedures are RECOMMENDED.

The LSP ingress is to send an MPLS echo request carrying the Extended BFD Discriminator TLV with the same BFD discriminator of the existing BFD session (one bootstrapped previously with the BFD Discriminator TLV), giving it an instance identifier. Once the transition of the existing BFD session is completed, then the LSP ingress can generate further MPLS echo request messages with the Extended BFD Discriminator TLV to bootstrap more BFD sessions.

7. Encapsulation

The encapsulation of BFD packets are the same as specified by [\[RFC5884\]](#)

8. Security Considerations

This document defines a mechanism to bootstrap multiple BFD sessions per FEC. 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.

9. IANA Considerations

9.1. Extended BFD Discriminator TLV

The IANA is requested to assign new value TBD1 for Extended BFD Discriminator TLV from the "Multiprotocol Label Switching Architecture (MPLS) Label Switched Paths (LSPs) Ping Parameters - TLVs" registry.

Value	Meaning	Reference
-----	-----	-----
TBD1	Extended BFD Discriminator TLV	this document

10. Acknowledgements

TBD

11. Contributing Authors

Girija Rao
Cisco Systems
Email: giraghav@cisco.com

Mallik Mudigonda
Cisco Systems
Email: mmudigon@cisco.com

12. 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.

Appendix A. Alternate format for the BFD Extended TLV

The BFD Extended TLV can be used to carry the Operation Type and the Operation Status (Op Status) bits that are defined below:

BFD Discriminator - When the Extended BFD Discriminator TLV is carried in the MPLS echo request, this field describes the BFD discriminator allocated for this BFD session by the LSP ingress. When the Extended BFD Discriminator TLV is carried in the MPLS echo reply, this field describes the BFD discriminator allocated for this BFD session by the LSP egress.

The Extended BFD Discriminator TLV in an MPLS echo request MUST have either Create(1) or Delete(2) operation type. The Extended BFD Discriminator TLV in an MPLS echo reply MUST have either CreateAck(3) or DeleteACK(4) operation type.

Authors' Addresses

Vengada Prasad Govindan
Cisco Systems

Email: venggovi@cisco.com

Nobo Akiya
Cisco Systems

Email: nobo@cisco.com

