

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 3, 2017

Z. Li
Sam Aldrin
Huawei Technologies
J. Tantsura
G. Mirsky
Ericsson
S. Zhuang
Huawei Technologies
October 30, 2016

BGP Link-State Extensions for Seamless BFD
draft-li-idr-bgp-ls-sbfd-extensions-00

Abstract

[RFC7880] defines a simplified mechanism to use Bidirectional Forwarding Detection (BFD) with large portions of negotiation aspects eliminated, thus providing benefits such as quick provisioning as well as improved control and flexibility to network nodes initiating the path monitoring. The link-state routing protocols (IS-IS, OSPF and OSPFv3) have been extended to advertise the Seamless BFD (S-BFD) Discriminators.

This draft defines extensions to the BGP Link-state address-family to carry the S-BFD Discriminators information via BGP.

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 May 3, 2017.

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.	Introduction	2
2.	Terminology	3
3.	Problem and Requirement	3
4.	BGP-LS Extensions for S-BFD Discriminators Exchanging	3
5.	Operations	6
6.	IANA Considerations	6
7.	Security Considerations	6
8.	Acknowledgements	6
9.	References	6
9.1.	Normative References	6
9.2.	Informative References	6
	Authors' Addresses	7

[1.](#) Introduction

[RFC7880] defines a simplified mechanism to use Bidirectional Forwarding Detection (BFD)[[RFC5880](#)] with large portions of negotiation aspects eliminated, thus providing benefits such as quick provisioning as well as improved control and flexibility to network nodes initiating the path monitoring.

[RFC7883] defines a mean of advertising one or more S-BFD Discriminators using the IS-IS Router Capability TLV. [[RFC7884](#)] defines a new OSPF Router Information (RI) TLV that allows OSPF routers to flood the S-BFD discriminator values associated with a target network identifier. This mechanism is applicable to both OSPFv2 and OSPFv3.

The link-state routing protocols (IS-IS, OSPF and OSPFv3) have been extended to advertise the S-BFD Discriminators. But flooding based propagation of the S-BFD Discriminators using IGP is limited by the perimeter of the IGP domain. For advertising the S-BFD Discriminators which span across IGP domains (e.g. multiple ASes), the Border Gateway Protocol (BGP) is better suited as its propagation perimeter is not limited like the IGPs.

This draft defines extensions to the BGP Link-state address-family to carry the S-BFD Discriminators information via BGP.

2. Terminology

This memo makes use of the terms defined in [[RFC7880](#)].

3. Problem and Requirement

Seamless MPLS [[I-D.ietf-mpls-seamless-mpls](#)] extends the core domain and integrates aggregation and access domains into a single MPLS domain. In a large network, the core and aggregation networks can be organized as different autonomous systems. Although the core and aggregation networks are segmented into different autonomous systems, but an E2E LSP will be created using hierarchical-labeled BGP LSPs based on iBGP-labeled unicast within each AS, and eBGP-labeled unicast to extend the LSP across AS boundaries. Meanwhile, the customer will see only two service-end points in the Seamless MPLS network. In order to detect the possible failure quickly and protect the network/trigger re-routing, BFD MAY be used for the Service Layer (e.g. for MPLS VPNs, PW) and the Transport Layer, so the need arises that the BFD session has to span across AS domain.

The link-state routing protocols (IS-IS, OSPF and OSPFv3) have been extended to advertise the S-BFD Discriminators. But flooding based propagation of the S-BFD Discriminators using IGPs is limited by the perimeter of the IGP domain. For advertising the S-BFD Discriminators which span across IGP domains (e.g. multiple ASes), the Border Gateway Protocol (BGP) is better suited as its propagation perimeter is not limited like the IGPs. This draft defines extensions requirement to the BGP Link-state address-family to carry the S-BFD Discriminators information via BGP.

4. BGP-LS Extensions for S-BFD Discriminators Exchanging

The BGP-LS NLRI can be a node NLRI, a link NLRI or a prefix NLRI. The corresponding BGP-LS attribute is a node attribute, a link attribute or a prefix attribute. BGP-LS [[RFC7752](#)] defines the TLVs that map link-state information to BGP-LS NLRI and BGP-LS attribute.

This document adds additional BGP-LS attribute TLVs to encode the S-BFD Discriminators information.

[RFC7880] defines the following TLVs to encode the S-BFD Discriminators information.

The ISIS Router CAPABILITY TLV as defined in [RFC4971] will be used to advertise S-BFD discriminators. A new Sub-TLV is defined as described below. S-BFD Discriminators Sub-TLV is formatted as specified in [RFC5305].

	No. of octets
+-----+	
Type (20)	1
+-----+	
Length (multiple of 4)	1
+-----+	
Discriminator Value(s)	4/Discriminator
:	:
+-----+	

Figure 1: S-BFD Discriminators Sub-TLV for ISIS

Inclusion of the S-BFD Discriminators sub-TLV in a Router Capability TLV is optional. Multiple S-BFD Discriminators sub-TLVs MAY be advertised by an IS.

[RFC7884] defines the following TLVs to encode the S-BFD Discriminators information. The format of the S-BFD Discriminator TLV is as follows:

0	1	2	3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1			
+-----+			
	Type		Length
+-----+			
	Discriminator 1		
+-----+			
	Discriminator 2 (Optional)		
+-----+			
	...		
+-----+			
	Discriminator n (Optional)		
+-----+			

Figure 2: S-BFD Discriminators TLV for OSPF


Type - S-BFD Discriminator TLV Type (11)

Length - Total length of the discriminator (Value field) in octets, not including the optional padding. The Length is a multiple of 4 octets, and consequently specifies how many Discriminators are included in the TLV.

Value - S-BFD network target discriminator value or values.

Routers that do not recognize the S-BFD Discriminator TLV Type MUST ignore the TLV. S-BFD discriminator is associated with the BFD Target Identifier type, which allows de-multiplexing to a specific task or service.

These TLVs are mapped to BGP-LS attribute TLVs in the following way. The new information in the Link-State NLRIs and attributes is encoded in Type/Length/Value triplets.



The 2 octet Type field values are defined in Table 1. The next 2 octet Length field encodes length of the rest of the TLV. The Value portion of the TLV is variable and is equal to the corresponding Value portion of the TLV defined in [\[RFC7883\]](#) and [\[RFC7884\]](#).

The following 'Node Attribute' TLVs are defined:

TLV Code	Description	Length	ISIS/OSPF
Point			TLV/Sub-TLV
TBD	S-BFD Discriminators	variable	TBD
...

Table 1: Node Attribute TLVs

These TLVs can ONLY be added to the Node Attribute associated with the Node NLRI that originates the corresponding S-BFD Discriminator TLV.

5. Operations

Existing BGP and BGP-LS operational procedures apply. No new operation procedures are defined in this document.

6. IANA Considerations

This document requests assigning code-points from the registry for BGP-LS attribute TLVs based on table Table 1.

7. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See [[RFC6952](#)] for details.

8. Acknowledgements

The authors would like to thank Nan Wu for his contributions to this work.

9. References

9.1. 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>>.

9.2. Informative References

- [I-D.ietf-mpls-seamless-mpls]
Leymann, N., Decraene, B., Filsfils, C., Konstantynowicz, M., and D. Steinberg, "Seamless MPLS Architecture", [draft-ietf-mpls-seamless-mpls-07](#) (work in progress), June 2014.
- [RFC4971] Vasseur, JP., Ed., Shen, N., Ed., and R. Aggarwal, Ed., "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", [RFC 4971](#), DOI 10.17487/RFC4971, July 2007, <<http://www.rfc-editor.org/info/rfc4971>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.

- [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>>.
- [RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", [RFC 6952](#), DOI 10.17487/RFC6952, May 2013, <<http://www.rfc-editor.org/info/rfc6952>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", [RFC 7752](#), DOI 10.17487/RFC7752, March 2016, <<http://www.rfc-editor.org/info/rfc7752>>.
- [RFC7880] Pignataro, C., Ward, D., Akiya, N., Bhatia, M., and S. Pallagatti, "Seamless Bidirectional Forwarding Detection (S-BFD)", [RFC 7880](#), DOI 10.17487/RFC7880, July 2016, <<http://www.rfc-editor.org/info/rfc7880>>.
- [RFC7883] Ginsberg, L., Akiya, N., and M. Chen, "Advertising Seamless Bidirectional Forwarding Detection (S-BFD) Discriminators in IS-IS", [RFC 7883](#), DOI 10.17487/RFC7883, July 2016, <<http://www.rfc-editor.org/info/rfc7883>>.
- [RFC7884] Pignataro, C., Bhatia, M., Aldrin, S., and T. Ranganath, "OSPF Extensions to Advertise Seamless Bidirectional Forwarding Detection (S-BFD) Target Discriminators", [RFC 7884](#), DOI 10.17487/RFC7884, July 2016, <<http://www.rfc-editor.org/info/rfc7884>>.

Authors' Addresses

Zhenbin Li
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Sam Aldrin
Huawei Technologies
2330 Central Expressway
Santa Clara CA 95051

Email: sam.aldrin@huawei.com

Jeff Tantsura
Ericsson
200 Holger Way
San Jose CA 95134
USA

Email: jeff.tantsura@ericsson.com

Greg Mirsky
Ericsson
300 Holger Way
San Jose CA 95134
USA

Email: gregory.mirsky@ericsson.com

Shunwan Zhuang
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: zhuangshunwan@huawei.com

