

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: August 26, 2019

Z. Li
Huawei
S. Aldrin
Google, Inc
J. Tantsura
Apstra
G. Mirsky
ZTE Corp.
S. Zhuang
Huawei
K. Talaulikar
Cisco Systems
February 22, 2019

BGP Link-State Extensions for Seamless BFD
draft-li-idr-bgp-ls-sbfd-extensions-03

Abstract

Seamless Bidirectional Forwarding Detection (S-BFD) 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 and OSPF) 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", "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.

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 August 26, 2019.

Copyright Notice

Copyright (c) 2019 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
2. Terminology	3
3. Problem and Requirement	3
4. BGP-LS Extensions for S-BFD Discriminator	4
5. IANA Considerations	6
6. Manageability Considerations	6
6.1. Operational Considerations	6
6.2. Management Considerations	6
7. Security Considerations	6
8. Acknowledgements	7
9. References	7
9.1. Normative References	7
9.2. Informative References	7
Authors' Addresses	8

1. Introduction

Seamless Bidirectional Forwarding Detection (S-BFD) [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.

For monitoring of a service path end-to-end via S-BFD, the headend/initiator node needs to know the S-BFD Discriminator of the destination/tail-end node of that service. The link-state routing protocols (IS-IS, OSPF and OSPFv3) have been extended to advertise the S-BFD Discriminators. With this a initiator node can learn the S-BFD discriminator for all nodes within its IGP area/level or optionally within the domain. With networks being divided into multiple IGP domains for scaling and operational considerations, the service endpoints that require end to end S-BFD monitoring often span across IGP domains.

BGP Link-State (BGP-LS) [RFC7752] enables the collection and distribution of IGP link-state topology information via BGP sessions across IGP areas/levels and domains. The S-BFD discriminator(s) of a node can thus be distributed along with the topology information via BGP-LS across IGP domains and even across multiple Autonomous Systems (AS) within an administrative domain.

This draft defines extensions to BGP-LS for carrying the S-BFD Discriminators information.

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 ASes. Although the core and aggregation networks are segmented into different ASes, an E2E LSP can be created using hierarchical BGP signaled LSPs based on iBGP labeled unicast within each AS, and eBGP labeled unicast to extend the LSP across AS boundaries. This provides a seamless MPLS transport connectivity for any two service end-points across the entire domain. In order to detect failures for such end to end services and trigger faster protection and/or re-routing, S-BFD MAY be used for the Service Layer (e.g. for MPLS VPNs, PW, etc.) or the Transport Layer monitoring. This brings up the need for setting up S-BFD session spanning across AS domains.

In a similar Segment Routing (SR) [RFC8402] multi-domain network, an end to end SR Policy [I-D.ietf-spring-segment-routing-policy] path may be provisioned between service end-points across domains either via local provisioning or by a controller or signalled from a Path Computation Engine (PCE). Monitoring using S-BFD can similarly be setup for such a SR Policy.

Extending the automatic discovery of S-BFD discriminators of nodes from within the IGP domain to across the administrative domain using BGP-LS enables setting up of S-BFD sessions on demand across IGP domains. The S-BFD discriminators for service end point nodes MAY be learnt by the PCE or a controller via the BGP-LS feed that it gets from across IGP domains and it can signal or provision the remote S-BFD discriminator on the initiator node on demand when S-BFD monitoring is required. The mechanisms for the signaling of the S-BFD discriminator from the PCE/controller to the initiator node and setup of the S-BFD session is outside the scope of this document.

Additionally, the service end-points themselves MAY also learn the S-BFD discriminator of the remote nodes themselves by receiving the BGP-LS feed via a route reflector (RR) or a centralized BGP Speaker that is consolidating the topology information across the domains. The initiator node can then itself setup the S-BFD session to the remote node without a controller/PCE assistance.

While this document takes examples of MPLS and SR paths, the S-BFD discriminator advertisement mechanism is applicable for any S-BFD use-case in general.

4. BGP-LS Extensions for S-BFD Discriminator

The BGP-LS [RFC7752] specifies the Node NLRI for advertisement of nodes and their attributes using the BGP-LS Attribute. The S-BFD discriminators of a node are considered as its node level attribute and advertised as such.

This document defines a new BGP-LS Attribute TLV called the S-BFD Discriminators TLV and its format is as follows:




Figure 1: S-BFD Discriminators TLV

where:

- o Type: TBD (see IANA Considerations Section 5)
- o Length: variable. Minimum of 8 octets and increments of 4 octets there on for each additional discriminator
- o Discriminators : multiples of 4 octets, each carrying a S-BFD local discriminator value of the node. At least one discriminator MUST be included in the TLV.

The S-BFD Discriminators TLV can only be added to the BGP-LS Attribute associated with the Node NLRI that originates the corresponding underlying IGP TLV/sub-TLV as described below. This information is derived from the protocol specific advertisements as below..

- o IS-IS, as defined by the S-BFD Discriminators sub-TLV in [RFC7883].
- o OSPFv2/OSPFv3, as defined by the S-BFD Discriminators TLV in [RFC7884].

When the node is not running any of the IGPs but running a protocol like BGP, then the locally provisioned S-BFD discriminators of the node MAY be originated as part of the BGP-LS attribute within the Node NLRI corresponding to the local node.

5. IANA Considerations

This document requests assigning code-points from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" based on table below. The column "IS-IS TLV/Sub-TLV" defined in the registry does not require any value and should be left empty.

Code Point	Description	Length
TBD	S-BFD Discriminators TLV	variable

6. Manageability Considerations

This section is structured as recommended in [RFC5706].

The new protocol extensions introduced in this document augment the existing IGP topology information that was distributed via [RFC7752]. Procedures and protocol extensions defined in this document do not affect the BGP protocol operations and management other than as discussed in the Manageability Considerations section of [RFC7752]. Specifically, the malformed NLRIs attribute tests in the Fault Management section of [RFC7752] now encompass the new TLVs for the BGP-LS NLRI in this document.

6.1. Operational Considerations

No additional operation considerations are defined in this document.

6.2. Management Considerations

No additional management considerations are defined in this document.

7. Security Considerations

The new protocol extensions introduced in this document augment the existing IGP topology information that was distributed via [RFC7752]. Procedures and protocol extensions defined in this document do not affect the BGP security model other than as discussed in the Security Considerations section of [RFC7752]. More specifically the aspects related to limiting the nodes and consumers with which the topology information is shared via BGP-LS to trusted entities within an administrative domain.

Advertising the S-BFD Discriminators via BGP-LS makes it possible for attackers to initiate S-BFD sessions using the advertised

information. The vulnerabilities this poses and how to mitigate them are discussed in [RFC7752].

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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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, <<https://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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc7884>>.
- [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>>.

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.
- [I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Sivabalan, S., daniel.voyer@bell.ca, d., bogdanov@google.com, b., and P. Mattes, "Segment Routing Policy Architecture", draft-ietf-spring-segment-routing-policy-02 (work in progress), October 2018.
- [RFC5706] Harrington, D., "Guidelines for Considering Operations and Management of New Protocols and Protocol Extensions", RFC 5706, DOI 10.17487/RFC5706, November 2009, <<https://www.rfc-editor.org/info/rfc5706>>.
- [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>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

Authors' Addresses

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

Email: lizhenbin@huawei.com

Sam Aldrin
Google, Inc

Email: aldrin.ietf@gmail.com

Jeff Tantsura
Apstra

Email: jefftant.ietf@gmail.com

Greg Mirsky
ZTE Corp.

Email: gregimirsky@gmail.com

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

Email: zhuangshunwan@huawei.com

Ketan Talaulikar
Cisco Systems
India

Email: ketant@cisco.com