

Internet Engineering Task Force (IETF)
Request for Comments: 9086
Category: Standards Track
ISSN: 2070-1721

S. Previdi
Huawei Technologies
K. Talaulikar, Ed.
C. Filsfils
Cisco Systems, Inc.
K. Patel
Arrcus, Inc.
S. Ray
Individual
J. Dong
Huawei Technologies
August 2021

Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment
Routing BGP Egress Peer Engineering

Abstract

A node steers a packet through a controlled set of instructions, called segments, by prepending the packet with a list of segment identifiers (SIDs). A segment can represent any instruction, topological or service based. SR segments allow steering a flow through any topological path and service chain while maintaining per-flow state only at the ingress node of the SR domain.

This document describes an extension to Border Gateway Protocol - Link State (BGP-LS) for advertisement of BGP Peering Segments along with their BGP peering node information so that efficient BGP Egress Peer Engineering (EPE) policies and strategies can be computed based on Segment Routing.

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 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9086>.

Copyright Notice

Copyright (c) 2021 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. Requirements Language
3. BGP Peering Segments
4. BGP-LS NLRI Advertisement for BGP Protocol

- 4.1. BGP Router-ID and Member AS Number
- 4.2. Mandatory BGP Node Descriptors
- 4.3. Optional BGP Node Descriptors
- 5. BGP-LS Attributes for BGP Peering Segments
 - 5.1. Advertisement of the PeerNode SID
 - 5.2. Advertisement of the PeerAdj SID
 - 5.3. Advertisement of the PeerSet SID
- 6. IANA Considerations
 - 6.1. New BGP-LS Protocol-ID
 - 6.2. Node Descriptors and Link Attribute TLVs
- 7. Manageability Considerations
- 8. Security Considerations
- 9. References
 - 9.1. Normative References
 - 9.2. Informative References

Acknowledgements
Contributors
Authors' Addresses

1. Introduction

Segment Routing (SR) leverages source routing. A node steers a packet through a controlled set of instructions, called segments, by prepending the packet with a list of segment identifiers (SIDs). A SID can represent any instruction, topological or service based. SR segments allows to enforce a flow through any topological path or service function while maintaining per-flow state only at the ingress node of the SR domain.

The SR architecture [RFC8402] defines three types of BGP Peering Segments that may be instantiated at a BGP node:

- * Peer Node Segment (PeerNode SID) : instruction to steer to a specific peer node
- * Peer Adjacency Segment (PeerAdj SID) : instruction to steer over a specific local interface towards a specific peer node
- * Peer Set Segment (PeerSet SID) : instruction to load-balance to a set of specific peer nodes

SR can be directly applied to either an MPLS data plane (SR-MPLS) with no change on the forwarding plane or to a modified IPv6 forwarding plane (SRv6).

This document describes extensions to the BGP - Link State Network Layer Reachability Information (BGP-LS NLRI) and the BGP-LS Attribute defined for BGP-LS [RFC7752] for advertising BGP peering segments from a BGP node along with its peering topology information (i.e., its peers, interfaces, and peering Autonomous Systems (ASes)) to enable computation of efficient BGP Egress Peer Engineering (BGP-EPE) policies and strategies using the SR-MPLS data plane. The corresponding extensions for SRv6 are specified in [BGPLS-SRV6].

[RFC9087] illustrates a centralized controller-based BGP Egress Peer Engineering solution involving SR path computation using the BGP Peering Segments. This use case comprises a centralized controller that learns the BGP Peering SIDs via BGP-LS and then uses this information to program a BGP-EPE policy at any node in the domain to perform traffic steering via a specific BGP egress node to specific External BGP (EBGP) peer(s) optionally also over a specific interface. The BGP-EPE policy can be realized using the SR Policy framework [SR-POLICY].

This document introduces a new BGP-LS Protocol-ID for BGP and defines new BGP-LS Node and Link Descriptor TLVs to facilitate advertising BGP-LS Link NLRI to represent the BGP peering topology. Further, it specifies the BGP-LS Attribute TLVs for advertisement of the BGP Peering Segments (i.e., PeerNode SID, PeerAdj SID, and PeerSet SID) to be advertised in the same BGP-LS Link NLRI.

2. 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.

3. BGP Peering Segments

As described in [RFC8402], a BGP-EPE-enabled Egress Provider Edge (PE) node instantiates SR Segments corresponding to its attached peers. These segments are called BGP Peering Segments or BGP Peering SIDs. In the case of EBGp, they enable the expression of source-routed interdomain paths.

An ingress border router of an AS may compose a list of SIDs to steer a flow along a selected path within the AS, towards a selected egress border router C of the AS, and to a specific EBGp peer. At minimum, a BGP-EPE policy applied at an ingress PE involves two SIDs: the Node SID of the chosen egress PE and then the BGP Peering SID for the chosen egress PE peer or peering interface.

Each BGP session MUST be described by a PeerNode SID. The description of the BGP session MAY be augmented by additional PeerAdj SIDs. Finally, multiple PeerNode SIDs or PeerAdj SIDs MAY be part of the same group/set in order to group EPE resources under a common PeerSet SID. These BGP Peering SIDs and their encoding are described in detail in Section 5.

The following BGP Peering SIDs need to be instantiated on a BGP router for each of its BGP peer sessions that are enabled for Egress Peer Engineering:

- * One PeerNode SID MUST be instantiated to describe the BGP peer session.
- * One or more PeerAdj SID MAY be instantiated corresponding to the underlying link(s) to the directly connected BGP peer session.
- * A PeerSet SID MAY be instantiated and additionally associated and shared between one or more PeerNode SIDs or PeerAdj SIDs.

While an egress point in a topology usually refers to EBGp sessions between external peers, there's nothing in the extensions defined in this document that would prevent the use of these extensions in the context of Internal BGP (IBGP) sessions. However, unlike EBGp sessions, which are generally between directly connected BGP routers also along the traffic forwarding path, IBGP peer sessions may be set up to BGP routers that are not in the forwarding path. As such, when the IBGP design includes sessions with route reflectors, a BGP router SHOULD NOT instantiate a BGP Peering SID for those sessions to peer nodes that are not in the forwarding path since the purpose of BGP Peering SID is to steer traffic to those specific peers. Thus, the applicability for IBGP peering may be limited to only those deployments where the IBGP peer is also along the forwarding data path.

Any BGP Peering SIDs instantiated on the node are advertised via BGP-LS Link NLRI type as described in the sections below. An illustration of the BGP Peering SIDs' allocations in a reference BGP peering topology along with the information carried in the BGP-LS Link NLRI and its corresponding BGP-LS Attribute are described in [RFC9087].

4. BGP-LS NLRI Advertisement for BGP Protocol

This section describes the BGP-LS NLRI encodings that describe the BGP peering and link connectivity between BGP routers.

This document specifies the advertisement of BGP peering topology

information via BGP-LS Link NLRI type, which requires use of a new BGP-LS Protocol-ID.

Protocol-ID	NLRI Information Source Protocol
7	BGP

Table 1: BGP-LS Protocol Identifier for BGP

The use of a new Protocol-ID allows separation and differentiation between the BGP-LS NLRIs carrying BGP information from the BGP-LS NLRIs carrying IGP link-state information defined in [RFC7752].

The BGP Peering information along with their Peering Segments are advertised using BGP-LS Link NLRI type with the Protocol-ID set to BGP. BGP-LS Link NLRI type uses the Descriptor TLVs and BGP-LS Attribute TLVs as defined in [RFC7752]. In order to correctly describe BGP nodes, new TLVs are defined in this section.

[RFC7752] defines BGP-LS Link NLRI type as follows:




Figure 1: BGP-LS Link NLRI

Node Descriptors and Link Descriptors are defined in [RFC7752].

4.1. BGP Router-ID and Member AS Number

Two new Node Descriptor TLVs are defined in this document:

- * BGP Router Identifier (BGP Router-ID):

Type: 516

Length: 4 octets

Value: 4-octet unsigned non-zero integer representing the BGP Identifier as defined in [RFC6286]

- * Member-AS Number (Member-ASN)

Type: 517

Length: 4 octets

Value: 4-octet unsigned non-zero integer representing the Member-AS Number [RFC5065]

4.2. Mandatory BGP Node Descriptors

The following Node Descriptor TLVs MUST be included in BGP-LS NLRI as Local Node Descriptors when distributing BGP information:

- * BGP Router-ID (TLV 516), which contains a valid BGP Identifier of the local BGP node.

- * Autonomous System Number (TLV 512) [RFC7752], which contains the Autonomous System Number (ASN) or AS Confederation Identifier (an ASN) [RFC5065], if confederations are used, of the local BGP node.

Note that Section 2.1 of [RFC6286] requires the BGP identifier (Router-ID) to be unique within an Autonomous System and non-zero. Therefore, the <ASN, BGP Router-ID> tuple is globally unique. Their use in the Node Descriptor helps map Link-State NLRIs with BGP protocol-ID to a unique BGP router in the administrative domain where BGP-LS is enabled.

The following Node Descriptor TLVs MUST be included in BGP-LS Link NLRI as Remote Node Descriptors when distributing BGP information:

- * BGP Router-ID (TLV 516), which contains the valid BGP Identifier of the peer BGP node.
- * Autonomous System Number (TLV 512) [RFC7752], which contains the ASN or the AS Confederation Identifier (an ASN) [RFC5065], if confederations are used, of the peer BGP node.

4.3. Optional BGP Node Descriptors

The following Node Descriptor TLVs MAY be included in BGP-LS NLRI as Local Node Descriptors when distributing BGP information:

- * Member-ASN (TLV 517), which contains the ASN of the confederation member (i.e., Member-AS Number), if BGP confederations are used, of the local BGP node.
- * Node Descriptors as defined in [RFC7752].

The following Node Descriptor TLVs MAY be included in BGP-LS Link NLRI as Remote Node Descriptors when distributing BGP information:

- * Member-ASN (TLV 517), which contains the ASN of the confederation member (i.e., Member-AS Number), if BGP confederations are used, of the peer BGP node.
- * Node Descriptors as defined in [RFC7752].

5. BGP-LS Attributes for BGP Peering Segments

This section defines the BGP-LS Attributes corresponding to the following BGP Peer Segment SIDs:

- * Peer Node Segment Identifier (PeerNode SID)
- * Peer Adjacency Segment Identifier (PeerAdj SID)
- * Peer Set Segment Identifier (PeerSet SID)

The following new BGP-LS Link Attribute TLVs are defined for use with BGP-LS Link NLRI for advertising BGP Peering SIDs:

TLV Code Point	Description
1101	PeerNode SID
1102	PeerAdj SID
1103	PeerSet SID

Table 2: BGP-LS TLV Code Points for BGP-EPE

PeerNode SID, PeerAdj SID, and PeerSet SID all have the same format

as defined below:




Figure 2: BGP Peering SIDs TLV Format

- * Type: 1101, 1102, or 1103 as listed in Table 2
- * Length: variable. Valid values are either 7 or 8 based on whether the encoding is done as a SID Index or a label.
- * Flags: one octet of flags with the following definition:




Figure 3: Peering SID TLV Flags Format

- V-Flag: Value Flag. If set, then the SID carries a label value. By default, the flag is SET.
 - L-Flag: Local Flag. If set, then the value/index carried by the SID has local significance. By default, the flag is SET.
 - B-Flag: Backup Flag. If set, the SID refers to a path that is eligible for protection using fast reroute (FRR). The computation of the backup forwarding path and its association with the BGP Peering SID forwarding entry is implementation specific. Section 3.6 of [RFC9087] discusses some of the possible ways of identifying backup paths for BGP Peering SIDs.
 - P-Flag: Persistent Flag: If set, the SID is persistently allocated, i.e., the SID value remains consistent across router restart and session/interface flap.
 - Rsvd bits: Reserved for future use and MUST be zero when originated and ignored when received.
- * Weight: 1 octet. The value represents the weight of the SID for the purpose of load balancing. An example use of the weight is described in [RFC8402].
 - * SID/Index/Label. According to the TLV length and the V- and L-Flag settings, it contains either:
 - A 3-octet local label where the 20 rightmost bits are used for encoding the label value. In this case, the V- and L-Flags MUST be SET.
 - A 4-octet index defining the offset in the Segment Routing Global Block (SRGB) [RFC8402] advertised by this router. In this case, the SRGB MUST be advertised using the extensions defined in [RFC9085].

The values of the PeerNode SID, PeerAdj SID, and PeerSet SID Sub-TLVs SHOULD be persistent across router restart.

When enabled for Egress Peer Engineering, the BGP router MUST include the PeerNode SID TLV in the BGP-LS Attribute for the BGP-LS Link NLRI corresponding to its BGP peering sessions. The PeerAdj SID and PeerSet SID TLVs MAY be included in the BGP-LS Attribute for the BGP-

LS Link NLRI.

Additional BGP-LS Link Attribute TLVs as defined in [RFC7752] MAY be included with the BGP-LS Link NLRI in order to advertise the characteristics of the peering link, e.g., one or more interface addresses (TLV 259 or TLV 261) of the underlying link(s) over which a multi-hop BGP peering session is set up may be included in the BGP-LS Attribute along with the PeerNode SID TLV.

5.1. Advertisement of the PeerNode SID

The PeerNode SID TLV includes a SID associated with the BGP peer node that is described by a BGP-LS Link NLRI as specified in Section 4.

The PeerNode SID, at the BGP node advertising it, has the following semantics (as defined in [RFC8402]):

- * SR operation: NEXT
- * Next-Hop: the connected peering node to which the segment is associated

The PeerNode SID is advertised with a BGP-LS Link NLRI, where:

- * Local Node Descriptors include:
 - Local BGP Router-ID (TLV 516) of the BGP-EPE-enabled Egress PE
 - Local ASN (TLV 512)
- * Remote Node Descriptors include:
 - Peer BGP Router-ID (TLV 516) (i.e., the peer BGP ID used in the BGP session)
 - Peer ASN (TLV 512)
- * Link Descriptors include the addresses used by the BGP session encoded using TLVs as defined in [RFC7752]:
 - IPv4 Interface Address (TLV 259) contains the BGP session IPv4 local address.
 - IPv4 Neighbor Address (TLV 260) contains the BGP session IPv4 peer address.
 - IPv6 Interface Address (TLV 261) contains the BGP session IPv6 local address.
 - IPv6 Neighbor Address (TLV 262) contains the BGP session IPv6 peer address.
- * Link Attribute TLVs include the PeerNode SID TLV as defined in Figure 2.

5.2. Advertisement of the PeerAdj SID

The PeerAdj SID TLV includes a SID associated with the underlying link to the BGP peer node that is described by a BGP-LS Link NLRI as specified in Section 4.

The PeerAdj SID, at the BGP node advertising it, has the following semantics (as defined in [RFC8402]):

- * SR operation: NEXT
- * Next-Hop: the interface peer address

The PeerAdj SID is advertised with a BGP-LS Link NLRI, where:

- * Local Node Descriptors include:

- Local BGP Router-ID (TLV 516) of the BGP-EPE-enabled Egress PE
- Local ASN (TLV 512)
- * Remote Node Descriptors include:
 - Peer BGP Router-ID (TLV 516) (i.e., the peer BGP ID used in the BGP session)
 - Peer ASN (TLV 512)
- * Link Descriptors MUST include the following TLV, as defined in [RFC7752]:
 - Link Local/Remote Identifiers (TLV 258) contains the 4-octet Link Local Identifier followed by the 4-octet Link Remote Identifier. The value 0 is used by default when the link remote identifier is unknown.
- * Additional Link Descriptors TLVs, as defined in [RFC7752], MAY also be included to describe the addresses corresponding to the link between the BGP routers:
 - IPv4 Interface Address (Sub-TLV 259) contains the address of the local interface through which the BGP session is established.
 - IPv6 Interface Address (Sub-TLV 261) contains the address of the local interface through which the BGP session is established.
 - IPv4 Neighbor Address (Sub-TLV 260) contains the IPv4 address of the peer interface used by the BGP session.
 - IPv6 Neighbor Address (Sub-TLV 262) contains the IPv6 address of the peer interface used by the BGP session.
- * Link Attribute TLVs include the PeerAdj SID TLV as defined in Figure 2.

5.3. Advertisement of the PeerSet SID

The PeerSet SID TLV includes a SID that is shared amongst BGP peer nodes or the underlying links that are described by BGP-LS Link NLRI as specified in Section 4.

The PeerSet SID, at the BGP node advertising it, has the following semantics (as defined in [RFC8402]):

- * SR operation: NEXT
- * Next-Hop: load-balance across any connected interface to any peer in the associated peer set

The PeerSet SID TLV containing the same SID value (encoded as defined in Figure 2) is included in the BGP-LS Attribute for all of the BGP-LS Link NLRI corresponding to the PeerNode or PeerAdj segments associated with the peer set.

6. IANA Considerations

This document defines:

- * A new Protocol-ID: BGP. The code point is from the "BGP-LS Protocol-IDs" registry.
- * Two new TLVs: BGP-Router-ID and BGP Confederation Member. The code points are in the "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry.

- * Three new BGP-LS Attribute TLVs: PeerNode SID, PeerAdj SID, and PeerSet SID. The code points are in the "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry.

6.1. New BGP-LS Protocol-ID

This document defines a new value in the registry "BGP-LS Protocol-IDs":

Protocol-ID	NLRI information source protocol	Reference
7	BGP	RFC 9086

Table 3: BGP-LS Protocol-ID

6.2. Node Descriptors and Link Attribute TLVs

This document defines five new TLVs in the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs":

- * Two new Node Descriptor TLVs
- * Three new Link Attribute TLVs

All five of the new code points are in the same registry: "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs".

The following new Node Descriptor TLVs are defined:

TLV Code Point	Description	Reference
516	BGP Router-ID	RFC 9086
517	BGP Confederation Member	RFC 9086

Table 4: BGP-LS Descriptor TLV Code Points

The following new Link Attribute TLVs are defined:

TLV Code Point	Description	Reference
1101	PeerNode SID	RFC 9086
1102	PeerAdj SID	RFC 9086
1103	PeerSet SID	RFC 9086

Table 5: BGP-LS Attribute TLV Code Points

7. Manageability Considerations

The new protocol extensions introduced in this document augment the existing IGP topology information BGP-LS distribution [RFC7752] by adding support for distribution of BGP peering topology information. As such, Section 6 of [RFC7752] (Manageability Considerations) applies to these new extensions as well.

Specifically, the malformed Link-State NLRI and BGP-LS Attribute tests for syntactic checks in Section 6.2.2 of [RFC7752] (Fault Management) now apply to the TLVs defined in this document. The semantic or content checking for the TLVs specified in this document and their association with the BGP-LS NLRI types or their associated BGP-LS Attributes is left to the consumer of the BGP-LS information (e.g., an application or a controller) and not the BGP protocol.

A consumer of the BGP-LS information retrieves this information from a BGP Speaker, over a BGP-LS session (refer to Sections 1 and 2 of [RFC7752]). The handling of semantic or content errors by the consumer would be dictated by the nature of its application usage and is hence beyond the scope of this document. It may be expected that an error detected in the NLRI Descriptor TLVs would result in that specific NLRI update being unusable and hence its update to be discarded along with an error log, whereas an error in Attribute TLVs would result in only that specific attribute being discarded with an error log.

The operator MUST be provided with the options of configuring, enabling, and disabling the advertisement of each of the PeerNode SID, PeerAdj SID, and PeerSet SID as well as control of which information is advertised to which internal or external peer. This is not different from what is required by a BGP speaker in terms of information origination and advertisement.

BGP Peering Segments are associated with the normal BGP routing peering sessions. However, the BGP peering information along with these Peering Segments themselves are advertised via a distinct BGP-LS peering session. It is expected that this isolation as described in [RFC7752] is followed when advertising BGP peering topology information via BGP-LS.

BGP-EPE functionality enables the capability for instantiation of an SR path for traffic engineering a flow via an egress BGP router to a specific peer, bypassing the normal BGP best-path routing for that flow and any routing policies implemented in BGP on that egress BGP router. As with any traffic-engineering solution, the controller or application implementing the policy needs to ensure that there is no looping or misrouting of traffic. Traffic counters corresponding to the MPLS label of the BGP Peering SID on the router would indicate the traffic being forwarded based on the specific EPE path. Monitoring these counters and the flows hitting the corresponding MPLS forwarding entry would help identify issues, if any, with traffic engineering over the EPE paths. Errors in the encoding or decoding of the SR information in the TLVs defined in this document may result in the unavailability of such information to a Centralized EPE Controller or incorrect information being made available to it. This may result in the controller not being able to perform the desired SR-based optimization functionality or performing it in an unexpected or inconsistent manner. The handling of such errors by applications like such a controller may be implementation specific and out of scope of this document.

8. Security Considerations

[RFC7752] defines BGP-LS NLRI to which the extensions defined in this document apply. Section 8 of [RFC7752] also applies to these extensions. The procedures and new TLVs defined in this document, by themselves, do not affect the BGP-LS security model discussed in [RFC7752].

BGP-EPE enables engineering of traffic when leaving the administrative domain via an egress BGP router. Therefore, precaution is necessary to ensure that the BGP peering information collected via BGP-LS is limited to specific consumers in a secure manner. Segment Routing operates within a trusted domain [RFC8402], and its security considerations also apply to BGP Peering Segments. The BGP-EPE policies are expected to be used entirely within this trusted SR domain (e.g., between multiple AS/domains within a single provider network).

The isolation of BGP-LS peering sessions is also required to ensure that BGP-LS topology information (including the newly added BGP peering topology) is not advertised to an external BGP peering session outside an administrative domain.

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>>.
- [RFC5065] Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, DOI 10.17487/RFC5065, August 2007, <<https://www.rfc-editor.org/info/rfc5065>>.
- [RFC6286] Chen, E. and J. Yuan, "Autonomous-System-Wide Unique BGP Identifier for BGP-4", RFC 6286, DOI 10.17487/RFC6286, June 2011, <<https://www.rfc-editor.org/info/rfc6286>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC9085] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", RFC 9085, DOI 10.17487/RFC9085, August 2021, <<https://www.rfc-editor.org/info/rfc9085>>.

9.2. Informative References

- [BGPLS-SRV6] Dawra, G., Filsfils, C., Talaulikar, K., Chen, M., Bernier, D., and B. Decraene, "BGP Link State Extensions for SRv6", Work in Progress, Internet-Draft, draft-ietf-idr-bgpls-srv6-ext-08, 8 June 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgpls-srv6-ext-08>>.
- [RFC9087] Filsfils, C., Ed., Previdi, S., Dawra, G., Ed., Aries, E., and D. Afanasiev, "Segment Routing Centralized BGP Egress Peer Engineering", RFC 9087, DOI 10.17487/RFC9087, August 2021, <<https://www.rfc-editor.org/info/rfc9087>>.
- [SR-POLICY] Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, draft-ietf-spring-segment-routing-policy-13, 28 May 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-segment-routing-policy-13>>.

Acknowledgements

The authors would like to thank Jakob Heitz, Howard Yang, Hannes Gredler, Peter Psenak, Arjun Sreekantiah, and Bruno Decraene for their feedback and comments. Susan Hares helped in improving the clarity of the document with her substantial contributions during her shepherd's review. The authors would also like to thank Alvaro Retana for his extensive review and comments, which helped correct issues and improve the document.

Contributors

Mach(Guoyi) Chen
Huawei Technologies
China

Email: mach.chen@huawei.com

Acee Lindem
Cisco Systems Inc.
United States of America

Email: acee@cisco.com

Authors' Addresses

Stefano Previdi
Huawei Technologies

Email: stefano@previdi.net

Ketan Talaulikar (editor)
Cisco Systems, Inc.
India

Email: ketant@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels
Belgium

Email: cfilsfil@cisco.com

Keyur Patel
Arrcus, Inc.

Email: Keyur@arrcus.com

Saikat Ray
Individual

Email: raysaikat@gmail.com

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing
100095
China

Email: jie.dong@huawei.com