
Stream: Internet Engineering Task Force (IETF)

RFC: [8669](#)

Category: Standards Track

Published: December 2019

ISSN: 2070-1721

Authors:

S. Previdi

Huawei Technologies

C. Filsfils

Cisco Systems

A. Lindem, Ed.

Cisco Systems

A. Sreekantiah

H. Gredler

RtBrick Inc.

RFC 8669

Segment Routing Prefix Segment Identifier Extensions for BGP

Abstract

Segment Routing (SR) leverages the source-routing paradigm. A node steers a packet through an ordered list of instructions called "segments". A segment can represent any instruction, topological or service based. The ingress node prepends an SR header to a packet containing a set of segment identifiers (SIDs). Each SID represents a topological or service-based instruction. Per-flow state is maintained only on the ingress node of the SR domain. An "SR domain" is defined as a single administrative domain for global SID assignment.

This document defines an optional, transitive BGP attribute for announcing information about BGP Prefix Segment Identifiers (BGP Prefix-SIDs) and the specification for SR-MPLS SIDs.

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/rfc8669>.

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. MPLS BGP Prefix-SID
- 3. BGP Prefix-SID Attribute
 - 3.1. Label-Index TLV
 - 3.2. Originator SRGB TLV
- 4. Receiving BGP Prefix-SID Attribute
 - 4.1. MPLS Data Plane: Labeled Unicast
- 5. Advertising BGP Prefix-SID Attribute
 - 5.1. MPLS Data Plane: Labeled Unicast
- 6. Error Handling of BGP Prefix-SID Attribute
- 7. IANA Considerations
- 8. Manageability Considerations
- 9. Security Considerations
- 10. References
 - 10.1. Normative References
 - 10.2. Informative References
- Acknowledgements
- Contributors
- Authors' Addresses

1. Introduction

The Segment Routing (SR) architecture leverages the source-routing paradigm. A segment represents either a topological instruction, such as "go to prefix P following shortest path", or a service instruction. Other types of segments may be defined in the future.

A segment is identified through a Segment Identifier (SID). An "SR domain" is defined as a single administrative domain for global SID assignment. It may be comprised of a single Autonomous System (AS) or multiple ASes under consolidated global SID administration. Typically, the ingress node of the SR domain prepends an SR header containing SIDs to an incoming packet.

As described in [RFC8402], when SR is applied to the MPLS data plane ([RFC8660]), the SID consists of a label.

[RFC8402] also describes how Segment Routing can be applied to an IPv6 data plane (SRv6) using an IPv6 routing header containing a stack of SR SIDs encoded as IPv6 addresses [IPv6-SRH]. The applicability and support for Segment Routing over IPv6 is beyond the scope of this document.

A BGP Prefix Segment is a BGP prefix with a Prefix-SID attached. A BGP Prefix-SID is always a global SID ([RFC8402]) within the SR domain and identifies an instruction to forward the packet over the Equal-Cost Multipath (ECMP) best path computed by BGP to the related prefix. The BGP Prefix-SID is the identifier of the BGP Prefix Segment. In this document, we always refer to the BGP Prefix Segment by the BGP Prefix-SID.

This document describes the BGP extensions to signal the BGP Prefix-SID. Specifically, this document defines a BGP attribute known as the "BGP Prefix-SID attribute" and specifies the rules to originate, receive, and handle error conditions for the attribute.

The BGP Prefix-SID attribute defined in this document can be attached to prefixes from Multiprotocol BGP IPv4/IPv6 Labeled Unicast ([RFC4760] [RFC8277]). Usage of the BGP Prefix-SID attribute for other Address Family Identifier (AFI) / Subsequent Address Family Identifier (SAFI) combinations is not defined herein but may be specified in future specifications.

[RFC8670] describes example use cases where the BGP Prefix-SID is used for the above AFI/SAFI combinations.

It should be noted that:

- A BGP Prefix-SID will be global across ASes when the interconnected ASes are part of the same SR domain. Alternatively, when interconnecting ASes, the ASBRs of each domain will have to handle the advertisement of unique SIDs. The mechanisms for such interconnection are outside the scope of the protocol extensions defined in this document.
- A BGP Prefix-SID **MAY** be attached to a BGP prefix. This implies that each prefix is advertised individually, reducing the ability to pack BGP advertisements (when sharing common attributes).

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.

2. MPLS BGP Prefix-SID

The BGP Prefix-SID is realized on the MPLS data plane ([RFC8660]) in the following way:

The operator assigns a globally unique label index, L_I , to a locally originated prefix of a BGP speaker N , which is advertised to all other BGP speakers in the SR domain.

According to [RFC8402], each BGP speaker is configured with a label block called the Segment Routing Global Block (SRGB). While [RFC8402] recommends using the same SRGB across all the nodes within the SR domain, the SRGB of a node is a local property and could be different on different speakers. The drawbacks of the use case where BGP speakers have different SRGBs are documented in [RFC8402] and [RFC8670].

If traffic engineering within the SR domain is required, each node may also be required to advertise topological information and Peer SIDs for each of its links and peers. This information is required to perform the explicit path computation and to express an explicit path as a list of SIDs. The advertisement of topological information and peer segments (Peer SIDs) is done through [BGPLS-SR-EPE].

If a prefix segment is to be included in an MPLS label stack, e.g., for traffic-engineering purposes, knowledge of the prefix originator's SRGB is required in order to compute the local label used by the originator.

This document assumes that Border Gateway Protocol - Link State (BGP-LS) is the preferred method for a collecting both peer segments (Peer SIDs) and SRGB information through [RFC7752], [BGPLS-SR-EPE], and [BGPLS-SR-EXT]. However, as an optional alternative for the advertisement of the local SRGB without the topology or the peer SIDs and, therefore, without applicability for TE, the Originator SRGB TLV of the BGP Prefix-SID attribute is specified in Section 3.2 of this document.

A BGP speaker will derive its local MPLS label L from the label index L_I and its local SRGB as described in [RFC8660]. The BGP speaker then programs the MPLS label L in its MPLS data plane as its incoming/local label for the prefix. See Section 4.1 for more details.

The outgoing label for the prefix is found in the Network Layer Reachability Information (NLRI) of the Multiprotocol BGP IPv4/IPv6 Labeled Unicast prefix advertisement as defined in [RFC8277]. The label index L_I is only used as a hint to derive the local/incoming label.

Section 3.1 of this document specifies the Label-Index TLV of the BGP Prefix-SID attribute; this TLV can be used to advertise the label index for a given prefix.

3. BGP Prefix-SID Attribute

The BGP Prefix-SID attribute is an optional, transitive BGP path attribute. The attribute type code 40 has been assigned by IANA (see Section 7).

The BGP Prefix-SID attribute is defined here to be a set of elements encoded as "Type/Length/Value" tuples (i.e., a set of TLVs). All BGP Prefix-SID attribute TLVs will start with a 1-octet type and a 2-octet length. The following TLVs are defined in this document:


- Label-Index TLV
- Originator SRGB TLV

The Label-Index and Originator SRGB TLVs are used only when SR is applied to the MPLS data plane.

For future extensibility, unknown TLVs **MUST** be ignored and propagated unmodified.

3.1. Label-Index TLV

The Label-Index TLV **MUST** be present in the BGP Prefix-SID attribute attached to IPv4/IPv6 Labeled Unicast prefixes ([RFC8277]). It **MUST** be ignored when received for other BGP AFI/SAFI combinations. The Label-Index TLV has the following format:



where:

Type: 1

Length: 7, the total length in octets of the value portion of the TLV.

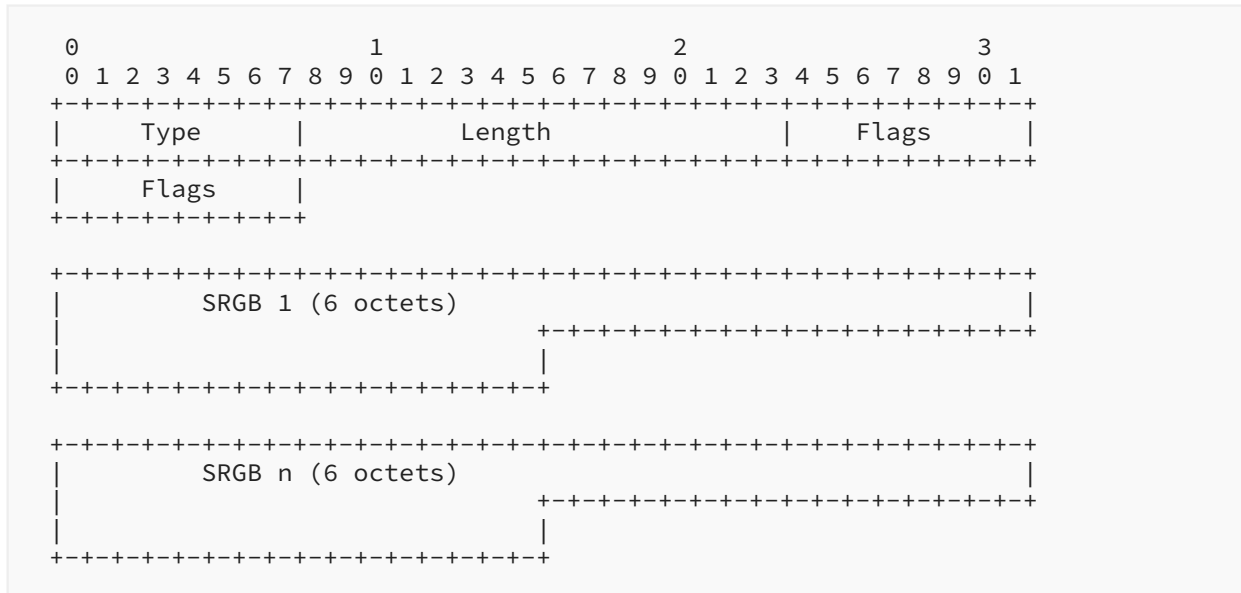
RESERVED: 8-bit field. It **MUST** be clear on transmission and **MUST** be ignored on reception.

Flags: 16 bits of flags. None are defined by this document. The Flags field **MUST** be clear on transmission and **MUST** be ignored on reception.

Label Index: 32-bit value representing the index value in the SRGB space.

3.2. Originator SRGB TLV

The Originator SRGB TLV is an optional TLV and has the following format:



where:

Type: 3

Length: The total length in octets of the value portion of the TLV: 2 + (non-zero multiple of 6).

Flags: 16 bits of flags. None are defined in this document. The Flags field **MUST** be clear on transmission and **MUST** be ignored on reception.

SRGB: 3 octets specifying the first label in the range followed by 3 octets specifying the number of labels in the range. Note that the SRGB field **MAY** appear multiple times. If the SRGB field appears multiple times, the SRGB consists of multiple ranges that are concatenated.

The Originator SRGB TLV contains the SRGB of the node originating the prefix to which the BGP Prefix-SID is attached. The Originator SRGB TLV **MUST NOT** be changed during the propagation of the BGP update. It is used to build SR policies when different SRGBs are used in the fabric, for example, [RFC8670].

Examples of how the receiving routers concatenate the ranges and build their neighbor's Segment Routing Global Block (SRGB) are included in [RFC8660].

The Originator SRGB TLV may only appear in a BGP Prefix-SID attribute attached to IPv4/IPv6 Labeled Unicast prefixes ([RFC8277]). It **MUST** be ignored when received for other BGP AFI/SAFI combinations. Since the Label-Index TLV is required for IPv4/IPv6 prefix applicability, the Originator SRGB TLV will be ignored if it is not specified in a manner consistent with Section 6.

If a BGP speaker receives a node's SRGB as an attribute of the BGP-LS Node NLRI and the BGP speaker also receives the same node's SRGB in a BGP Prefix-SID attribute, then the received values should be the same. If the values are different, the values advertised in the BGP-LS NLRI **SHOULD** be preferred, and an error should be logged.

4. Receiving BGP Prefix-SID Attribute

A BGP speaker receiving a BGP Prefix-SID attribute from an External BGP (EBGP) neighbor residing outside the boundaries of the SR domain **MUST** discard the attribute unless it is configured to accept the attribute from the EBGP neighbor. A BGP speaker **SHOULD** log an error for further analysis when discarding an attribute.

4.1. MPLS Data Plane: Labeled Unicast

A BGP session supporting the Multiprotocol BGP Labeled IPv4 or IPv6 Unicast ([RFC8277]) AFI/SAFI is required.

When the BGP Prefix-SID attribute is attached to a BGP Labeled IPv4 or IPv6 Unicast [RFC8277] AFI/SAFI, it **MUST** contain the Label-Index TLV and **MAY** contain the Originator SRGB TLV. A BGP Prefix-SID attribute received without a Label-Index TLV **MUST** be considered to be "invalid" by the receiving speaker.

The label index provides guidance to the receiving BGP speaker as to the incoming label that **SHOULD** be allocated to the prefix.

A BGP speaker may be locally configured with an SRGB=[SRGB_Start, SRGB_End]. The preferred method for deriving the SRGB is a matter of local node configuration.

The mechanisms through which a given label-index value is assigned to a given prefix are outside the scope of this document.

Given a label index L_I , we refer to $(L = L_I + \text{SRGB_Start})$ as the derived label. A BGP Prefix-SID attribute is designated "conflicting" for a speaker M if the derived label value L lies outside the SRGB configured on M. Otherwise, the Label-Index TLV is designated "acceptable" to speaker M.

If multiple different prefixes are received with the same label index, all of the different prefixes **MUST** have their BGP Prefix-SID attribute considered to be "conflicting".

If multiple valid paths for the same prefix are received from multiple BGP speakers or, in the case of [RFC7911], from the same BGP speaker, and the BGP Prefix-SID attributes do not contain the same label index, then the label index from the best path BGP Prefix-SID attribute **SHOULD** be chosen with a notable exception being when [RFC5004] is being used to dampen route changes.

When a BGP speaker receives a path from a neighbor with an "acceptable" BGP Prefix-SID attribute and that path is selected as the best path, it **SHOULD** program the derived label as the label for the prefix in its local MPLS data plane.

When a BGP speaker receives a path from a neighbor with an "invalid" or "conflicting" BGP Prefix-SID attribute, or when a BGP speaker receives a path from a neighbor with a BGP Prefix-SID attribute but is unable to process it (e.g., local policy disables the functionality), it **MUST** ignore the BGP Prefix-SID attribute. For the purposes of label allocation, a BGP speaker **MUST** assign a local (also called dynamic) label (non-SRGB) for such a prefix as per classic Multiprotocol BGP IPv4/IPv6 Labeled Unicast ([RFC8277]) operation.

In the case of an "invalid" BGP Prefix-SID attribute, a BGP speaker **MUST** follow the error-handling rules specified in Section 6. A BGP speaker **SHOULD** log an error for further analysis. In the case of a "conflicting" BGP Prefix-SID attribute, a BGP speaker **SHOULD NOT** treat it as an error and **SHOULD** propagate the attribute unchanged. A BGP speaker **SHOULD** log a warning for further analysis, i.e., in the case the conflict is not due to a label-index transition.

When a BGP Prefix-SID attribute changes and transitions from "conflicting" to "acceptable", the BGP Prefix-SID attributes for other prefixes may also transition to "acceptable" as well. Implementations **SHOULD** ensure all impacted prefixes revert to using the label indices corresponding to these newly "acceptable" BGP Prefix-SID attributes.

The outgoing label is always programmed as per classic Multiprotocol BGP IPv4/IPv6 Labeled Unicast ([RFC8277]) operation. Specifically, a BGP speaker receiving a prefix with a BGP Prefix-SID attribute and a label NLRI field of Implicit NULL [RFC3032] from a neighbor **MUST** adhere to standard behavior and program its MPLS data plane to pop the top label when forwarding traffic to the prefix. The label NLRI defines the outbound label that **MUST** be used by the receiving node.

5. Advertising BGP Prefix-SID Attribute

The BGP Prefix-SID attribute **MAY** be attached to BGP IPv4/IPv6 Labeled Unicast prefixes [RFC8277]. In order to prevent distribution of the BGP Prefix-SID attribute beyond its intended scope of applicability, attribute filtering **SHOULD** be deployed to remove the BGP Prefix-SID attribute at the administrative boundary of the SR domain.

A BGP speaker that advertises a path received from one of its neighbors **SHOULD** advertise the BGP Prefix-SID received with the path without modification as long as the BGP Prefix-SID was acceptable. If the path did not come with a BGP Prefix-SID attribute, the speaker **MAY** attach a BGP Prefix-SID to the path if configured to do so. The content of the TLVs present in the BGP Prefix-SID is determined by the configuration.

5.1. MPLS Data Plane: Labeled Unicast

A BGP speaker that originates a prefix attaches the BGP Prefix-SID attribute when it advertises the prefix to its neighbors via Multiprotocol BGP IPv4/IPv6 Labeled Unicast ([RFC8277]). The value of the label index in the Label-Index TLV is determined by configuration.

A BGP speaker that originates a BGP Prefix-SID attribute **MAY** optionally announce the Originator SRGB TLV along with the mandatory Label-Index TLV. The content of the Originator SRGB TLV is determined by configuration.

Since the label-index value must be unique within an SR domain, by default an implementation **SHOULD NOT** advertise the BGP Prefix-SID attribute outside an AS unless it is explicitly configured to do so.

In all cases, the Label field of the advertised NLRI ([RFC8277] [RFC4364]) **MUST** be set to the local/incoming label programmed in the MPLS data plane for the given advertised prefix. If the prefix is associated with one of the BGP speaker's interfaces, this is the usual MPLS label (such as the Implicit or Explicit NULL label [RFC3032]).

6. Error Handling of BGP Prefix-SID Attribute

When a BGP speaker receives a BGP UPDATE message containing a malformed or invalid BGP Prefix-SID attribute attached to an IPv4/IPv6 Labeled Unicast prefix ([RFC8277]), it **MUST** ignore the received BGP Prefix-SID attribute and not advertise it to other BGP peers. In this context, a malformed BGP Prefix-SID attribute is one that cannot be parsed due to not meeting the minimum attribute length requirement, containing a TLV length that doesn't conform to the length constraints for the TLV, or containing a TLV length that would extend beyond the end of the attribute (as defined by the attribute length). This is equivalent to the "Attribute discard" action specified in [RFC7606]. When discarding an attribute, a BGP speaker **SHOULD** log an error for further analysis.

As per [RFC7606], if the BGP Prefix-SID attribute appears more than once in an UPDATE message, all the occurrences of the attribute other than the first one **SHALL** be discarded and the UPDATE message will continue to be processed. Similarly, if a recognized TLV appears more than once in a BGP Prefix-SID attribute while the specification only allows for a single occurrence, then all the occurrences of the TLV other than the first one **SHALL** be discarded and the Prefix-SID attribute will continue to be processed.

For future extensibility, unknown TLVs **MUST** be ignored and propagated unmodified.

7. IANA Considerations

This document defines a BGP path attribute known as the BGP Prefix-SID attribute. IANA has assigned attribute code type 40 to the BGP Prefix-SID attribute from the "BGP Path Attributes" registry.

This document defines two TLVs for the BGP Prefix-SID attribute. These TLVs have been registered with IANA. IANA has created a registry for BGP Prefix-SID Attribute TLVs as follows:

Under the "Border Gateway Protocol (BGP) Parameters" registry, the new registry titled "BGP Prefix-SID TLV Types" has been created and points to this document as the reference.

Registration Procedure(s):

Values 1-254, Expert Review as defined in [RFC8126]

Values 0 and 255, Reserved

Value	Type	Reference
0	Reserved	This document
1	Label-Index	This document
2	Deprecated	This document
3	Originator SRGB	This document
4-254	Unassigned	
255	Reserved	This document

Table 1: BGP Prefix-SID TLV Types

The value 2 previously corresponded to the IPv6 SID TLV, which was specified in previous versions of this document. It was removed, and use of the BGP Prefix-SID for Segment Routing over the IPv6 data plane [RFC8402] has been deferred to future specifications.

IANA has also created the "BGP Prefix-SID Label-Index TLV Flags" registry under the "Border Gateway Protocol (BGP) Parameters" registry, with a reference to this document. Initially, this 16-bit flags registry is empty. The registration policy for flag bits is Expert Review [RFC8126], consistent with the "BGP Prefix-SID TLV Types" registry.

Finally, IANA has created the "BGP Prefix-SID Originator SRGB TLV Flags" registry under the "Border Gateway Protocol (BGP) Parameters" registry, with a reference to this document. Initially, this 16-bit flags registry is empty. The registration policy for flag bits is Expert Review [RFC8126] consistent with the BGP Prefix-SID TLV Types registry.

The designated experts must be good and faithful stewards of the above registries, ensuring that each request is legitimate and corresponds to a viable use case. Given the limited number of bits in the flags registries and the applicability to a single TLV, additional scrutiny should be afforded to requests for flag-bit allocation. In general, no single use case should require more than one flag bit and, should the use case require more, alternate encodings using new TLVs should be considered.

8. Manageability Considerations

This document defines a BGP attribute to address use cases such as the one described in [RFC8670]. It is assumed that advertisement of the BGP Prefix-SID attribute is controlled by the operator in order to:

- Prevent undesired origination/advertisement of the BGP Prefix-SID attribute. By default, a BGP Prefix-SID attribute **SHOULD NOT** be attached to a prefix and advertised. Hence, BGP Prefix-SID Advertisement **SHOULD** require explicit enablement.

- Prevent any undesired propagation of the BGP Prefix-SID attribute. By default, the BGP Prefix-SID is not advertised outside the boundary of a single SR/administrative domain that may include one or more ASes. The propagation to other ASes **MUST** be explicitly configured.

The deployment model described in [RFC8670] assumes multiple ASes under a common administrative domain. For this use case, the BGP Prefix-SID Advertisement is applicable to the inter-AS context, i.e., EBGp, while it is confined to a single administrative domain.

9. Security Considerations

This document introduces a BGP attribute (BGP Prefix-SID), which inherits the security considerations expressed in: [RFC4271], [RFC8277], and [RFC8402].

When advertised using BGPsec as described in [RFC8205], the BGP Prefix-SID attribute doesn't impose any unique security considerations. It should be noted that the BGP Prefix-SID attribute is not protected by the BGPsec signatures.

It should be noted that, as described in Section 8, this document refers to a deployment model where all nodes are under the single administrative domain. In this context, we assume that the operator doesn't want to leak any information related to internal prefixes and topology outside of the administrative domain. The internal information includes the BGP Prefix-SID. In order to prevent such leaking, the common BGP mechanisms (filters) are applied at the boundary of the SR/administrative domain. Local BGP-attribute-filtering policies and mechanisms are not standardized and, consequently, are beyond the scope of this document.

To prevent a Denial-of-Service (DoS) or Distributed-Denial-of-Service (DDoS) attack due to excessive BGP updates with an invalid or conflicting BGP Prefix-SID attribute, error log message rate limiting as well as suppression of duplicate error log messages **SHOULD** be deployed.

Since BGP-LS is the preferred method for advertising SRGB information, the BGP speaker **SHOULD** log an error if a BGP Prefix-SID attribute is received with SRGB information different from that received as an attribute of the same node's BGP-LS Node NLRI.

10. References

10.1. 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>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4364]

- Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.
- [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <<https://www.rfc-editor.org/info/rfc7911>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.
- [RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <<https://www.rfc-editor.org/info/rfc8205>>.
- [RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017, <<https://www.rfc-editor.org/info/rfc8277>>.
- [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>>.
- [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.

10.2. Informative References

- [BGPLS-SR-EPE] Previdi, S., Talaulikar, K., Filsfils, C., Patel, K., Ray, S., and J. Dong, "BGP-LS extensions for Segment Routing BGP Egress Peer Engineering", Work in Progress, Internet-Draft, draft-ietf-idr-bgpls-segment-routing-epe-19, 16 May 2019, <<https://tools.ietf.org/html/draft-ietf-idr-bgpls-segment-routing-epe-19>>.
- [BGPLS-SR-EXT] Previdi, S., Talaulikar, K., Filsfils, C., Gredler, H., and M. Chen, "BGP Link-State extensions for Segment Routing", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-segment-routing-ext-16, 27 June 2019, <<https://tools.ietf.org/html/draft-ietf-idr-bgp-ls-segment-routing-ext-16>>.

- [IPv6-SRH]** Filsfils, C., Dukes, D., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", Work in Progress, Internet-Draft, draft-ietf-6man-segment-routing-header-26, 22 October 2019, <<https://tools.ietf.org/html/draft-ietf-6man-segment-routing-header-26>>.
- [RFC3032]** Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.
- [RFC5004]** Chen, E. and S. Sangli, "Avoid BGP Best Path Transitions from One External to Another", RFC 5004, DOI 10.17487/RFC5004, September 2007, <<https://www.rfc-editor.org/info/rfc5004>>.
- [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>>.
- [RFC8670]** Filsfils, C., Ed., Previdi, S., Dawra, G., Aries, E., and P. Lapukhov, "BGP Prefix Segment in Large-Scale Data Centers", RFC 8670, DOI 10.17487/RFC8670, December 2019, <<https://www.rfc-editor.org/info/rfc8670>>.

Acknowledgements

The authors would like to thank Satya Mohanty for his contribution to this document.

The authors would like to thank Alvaro Retana for substantive comments as part of the Routing AD review.

The authors would like to thank Bruno Decraene for substantive comments and suggested text as part of the Routing Directorate review.

The authors would like to thank Shyam Sethuram for comments and discussion of TLV processing and validation.

The authors would like to thank Robert Raszuk for comments and suggestions regarding the MPLS data-plane behavior.

The authors would like to thank Krishna Deevi, Juan Alcaide, Howard Yang, and Jakob Heitz for discussions on conflicting BGP Prefix-SID label indices and BGP add paths.

The authors would like to thank Peter Yee, Tony Przygienda, Mirja Kuhlewind, Alexey Melnikov, Eric Rescorla, Suresh Krishnan, Warren Kumari, Ben Campbell Sue Hares, and Martin Vigoureur for IDR Working Group last call, IETF Last Call, directorate, and IESG reviews.

Contributors

Keyur Patel
Arrcus, Inc.
United States of America

Email: Keyur@arrcus.com

Saikat Ray
Unaffiliated
United States of America

Email: raysaikat@gmail.com

Authors' Addresses

Stefano Previdi

Huawei Technologies
Italy

Email: stefano@previdi.net

Clarence Filsfil

Cisco Systems
Brussels
Belgium

Email: cfilsfil@cisco.com

Acee Lindem (EDITOR)

Cisco Systems
301 Midenhall Way
Cary, NC, 27513
United States of America

Email: acee@cisco.com

Arjun Sreekantiah

Email: arjunhrs@gmail.com

Hannes Gredler

RtBrick Inc.

Email: hannes@rtbrick.com