

Internet Engineering Task Force (IETF)
Request for Comments: 9830
Updates: 9012
Category: Standards Track
ISSN: 2070-1721

S. Previdi
Huawei Technologies
C. Filsfils
K. Talaulikar, Ed.
Cisco Systems
P. Mattes
Microsoft
D. Jain
Google
September 2025

Advertising Segment Routing Policies in BGP

Abstract

A Segment Routing (SR) Policy is an ordered list of segments (also referred to as "instructions") that define a source-routed policy. An SR Policy consists of one or more Candidate Paths (CPs), each comprising one or more segment lists. A headend can be provisioned with these CPs using various mechanisms such as Command-Line Interface (CLI), Network Configuration Protocol (NETCONF), Path Computation Element Communication Protocol (PCEP), or BGP.

This document specifies how BGP can be used to distribute SR Policy CPs. It introduces a BGP SAFI for advertising a CP of an SR Policy and defines sub-TLVs for the Tunnel Encapsulation Attribute to signal information related to these CPs.

Furthermore, this document updates RFC 9012 by extending the Color Extended Community to support additional steering modes over SR Policy.

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/rfc9830>.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction
1.1.	Requirements Language
2.	SR Policy Encoding
2.1.	SR Policy SAFI and NLRI
2.2.	SR Policy and Tunnel Encapsulation Attribute
2.3.	Applicability of Tunnel Encapsulation Attribute Sub-TLVs
2.4.	SR Policy Sub-TLVs
2.4.1.	Preference Sub-TLV
2.4.2.	Binding SID Sub-TLV
2.4.3.	SRv6 Binding SID Sub-TLV
2.4.4.	Segment List Sub-TLV
2.4.5.	Explicit NULL Label Policy Sub-TLV
2.4.6.	SR Policy Priority Sub-TLV
2.4.7.	SR Policy Candidate Path Name Sub-TLV
2.4.8.	SR Policy Name Sub-TLV
3.	Color Extended Community
4.	SR Policy Operations
4.1.	Advertisement of SR Policies
4.2.	Reception of an SR Policy NLRI
4.2.1.	Validation of an SR Policy NLRI
4.2.2.	Eligibility for Local Use of an SR Policy NLRI
4.2.3.	Propagation of an SR Policy
5.	Error Handling and Fault Management
6.	IANA Considerations
6.1.	Subsequent Address Family Identifiers (SAFI) Parameters
6.2.	BGP Tunnel Encapsulation Attribute Tunnel Types
6.3.	BGP Tunnel Encapsulation Attribute Sub-TLVs
6.4.	Color Extended Community Flags
6.5.	SR Policy Segment List Sub-TLVs
6.6.	SR Policy Binding SID Flags
6.7.	SR Policy SRv6 Binding SID Flags
6.8.	SR Policy Segment Flags
6.9.	Color Extended Community Color-Only Types
6.10.	SR Policy ENLP Values
7.	Security Considerations
8.	Manageability Considerations
9.	References
9.1.	Normative References
9.2.	Informative References
	Acknowledgments
	Contributors
	Authors' Addresses

1. Introduction

Segment Routing (SR) [RFC8402] allows a headend node to steer a packet flow along a specific path. Intermediate per-path states are eliminated thanks to source routing.

The headend node is said to steer a flow into an SR Policy [RFC9256].

The packets steered into an SR Policy carry an ordered list of segments associated with that SR Policy.

[RFC9256] further details the concepts of SR Policy and steering into an SR Policy. These apply equally to the SR-MPLS and Segment Routing over IPv6 (SRv6) data plane instantiations of Segment Routing using SR-MPLS and SRv6 Segment Identifiers (SIDs) as described in [RFC8402]. [RFC8660] describes the representation and processing of this ordered list of segments as an MPLS label stack for SR-MPLS. [RFC8754] and [RFC8986] describe the same for SRv6 with the use of the Segment Routing Header (SRH).

The functionality related to SR Policy described in [RFC9256] can be

conceptually viewed as being incorporated in an SR Policy Module (SRPM). The following is a reminder of the high-level functionality of SRPM:

- * Learning multiple CPs for an SR Policy via various mechanisms (CLI, NETCONF, PCEP, or BGP).
- * Selection of the best CP for an SR Policy.
- * Associating a Binding SID (BSID) to the selected CP of an SR Policy.
- * Installation of the selected CP and its BSID in the forwarding plane.

This document specifies the use of BGP to distribute one or more of the CPs of an SR Policy to the headend of that SR Policy. The document describes the functionality provided by BGP and, as appropriate, provides references for the functionality, which is outside the scope of BGP (i.e., resides within SRPM on the headend node).

This document specifies a way of representing SR Policy CPs in BGP UPDATE messages. BGP can then be used to propagate the SR Policy CPs to the headend nodes in a network. The usual BGP rules for BGP propagation and best-path selection are used. At the headend of a specific SR Policy, this will result in one or more CPs being installed into the "BGP table". These paths are then passed to the SRPM. The SRPM may compare them to CPs learned via other mechanisms and will choose one or more paths to be installed in the data plane. BGP itself does not install SR Policy CPs into the data plane.

This document introduces a BGP Subsequent Address Family Identifier (SAFI) for IPv4 and IPv6 address families. In BGP UPDATE messages of those AFI/SAFIs, the Network Layer Reachability Information (NLRI) identifies an SR Policy CP while the attributes encode the segment lists and other details of that SR Policy CP.

While, for simplicity, the text in this document states that BGP advertises an SR Policy, it is to be understood that BGP advertises a CP of an SR Policy and that this SR Policy might have several other CPs provided via BGP (via an NLRI with a different distinguisher as defined in Section 2.1), PCEP, NETCONF, or local policy configuration.

Typically, an SR Policy Controller [RFC9256] defines the set of policies and advertises them to SR Policy headend routers (typically ingress routers). These SR Policy advertisements use the BGP extensions defined in this document. In most cases, the SR Policy advertisement is tailored for a specific SR Policy headend; consequently, it may be transmitted over a direct BGP session (i.e., without intermediate BGP hops) to that headend and is not propagated any further. In such cases, the SR Policy advertisements will not traverse any Route Reflector (RR) (see [RFC4456] and Section 4.2.3).

Alternatively, a BGP egress router may advertise SR Policies that represent paths that terminate on it. In such cases, the router can send these policies directly to each headend over a dedicated BGP session, without necessitating any further propagation of the SR Policy.

In some situations, it is undesirable for a controller or BGP egress router to have a BGP session to each SR Policy headend. In these situations, BGP RRs may be used to propagate the advertisements. In certain other deployments, it may be necessary for the advertisement to propagate through a sequence of one or more Autonomous Systems

(ASes) within an SR Domain (refer to Section 7 for the associated security considerations). To make this possible, an attribute needs to be attached to the advertisement that enables a BGP speaker to determine whether it is intended to be a headend for the advertised SR Policy. This is done by attaching one or more Route Target extended communities to the advertisement [RFC4360].

The BGP extensions for the advertisement of SR Policies include following components:

- * A SAFI whose NLRIs identify an SR Policy CP.
- * A Tunnel Type identifier for SR Policy and a set of sub-TLVs to be inserted into the Tunnel Encapsulation Attribute (as defined in [RFC9012]) specifying segment lists of the SR Policy CP as well as other information about the SR Policy.
- * One or more IPv4 address-specific format Route Target extended community ([RFC4360]) attached to the SR Policy CP advertisement that indicates the intended headend of such an SR Policy CP advertisement.

The SR Policy SAFI route updates utilize the Tunnel Encapsulation Attribute to signal an SR Policy, which itself functions as a tunnel. This usage differs notably from the approach described in [RFC9012], where the Tunnel Encapsulation Attribute is associated with a BGP route update (e.g., for Internet or VPN routes) to specify the tunnel used for forwarding traffic. This document does not modify or supersede the usage of the Tunnel Encapsulation Attribute for existing AFI/SAFIs as defined in [RFC9012]. Details regarding the processing of the Tunnel Encapsulation Attribute for the SR Policy SAFI are provided in Sections 2.2 and 2.3.

The northbound advertisement of the operational state of the SR Policy CPs as part of BGP - Link State (BGP-LS) [RFC9552] topology information is specified in [BGP-LS-SR-POLICY].

The signaling of Dynamic and Composite CPs (Sections 5.2 and 5.3, respectively, of [RFC9256]) is outside the scope of this document.

The Color Extended Community (as defined in [RFC9012]) is used to steer traffic into an SR Policy, as described in Section 8.8 of [RFC9256]. Section 3 of this document updates [RFC9012] with modifications to the format of the Flags field of the Color Extended Community by using the two leftmost bits of that field.

1.1. 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.

2. SR Policy Encoding

2.1. SR Policy SAFI and NLRI

The SR Policy SAFI with code point 73 is introduced in this document. The AFI used MUST be IPv4(1) or IPv6(2).

The SR Policy SAFI uses the NLRI format defined as follows:

```
+-----+
| NLRI Length      | 1 octet
+-----+
```

Distinguisher	4 octets
+-----+	
Color	4 octets
+-----+	
Endpoint	4 or 16 octets
+-----+	

Figure 1: SR Policy SAFI Format

Where:

NLRI Length: 1 octet indicating the length expressed in bits as defined in [RFC4760]. When AFI = 1, the value MUST be 96; when AFI = 2, the value MUST be 192.

Distinguisher: 4-octet value uniquely identifying the SR Policy in the context of <Color, Endpoint> tuple. The distinguisher has no semantic value. It is used by the SR Policy originator to form unique NLRIs the following situations:

- * to differentiate multiple CPs of the same SR Policy
- * to differentiate CPs meant for different headends but having the same Color and Endpoint

The distinguisher is the discriminator of the SR Policy CP as specified in Section 2.5 of [RFC9256].

Color: 4 octets that carry an unsigned non-zero integer value indicating the Color of the SR Policy as specified in Section 2.1 of [RFC9256]. The Color is used to match the Color of the destination prefixes to steer traffic into the SR Policy as specified in Section 8 of [RFC9256].

Endpoint: a value that identifies the Endpoint of an SR Policy. The Endpoint may represent a single node or a set of nodes (e.g., an anycast address). The Endpoint is an IPv4 (4-octet) address or an IPv6 (16-octet) address according to the AFI of the NLRI. The address can be either unicast or an unspecified address (0.0.0.0 for IPv4, :: for IPv6), known as a null Endpoint as specified in Section 2.1 of [RFC9256].

The Color and Endpoint are used to automate the steering of BGP service routes on an SR Policy as described in Section 8 of [RFC9256].

The NLRI containing an SR Policy CP is carried in a BGP UPDATE message [RFC4271] using BGP multiprotocol extensions [RFC4760] with an AFI of 1 or 2 (IPv4 or IPv6) and with a SAFI of 73. The fault management and error handling in the encoding of the NLRI are specified in Section 5.

A BGP UPDATE message that carries the MP_REACH_NLRI or MP_UNREACH_NLRI attribute with the SR Policy SAFI MUST also carry the BGP mandatory attributes. In addition, the BGP UPDATE message MAY also contain any of the BGP optional attributes.

The next-hop network address field in SR Policy SAFI (73) updates may be either a 4-octet IPv4 address or a 16-octet IPv6 address, independent of the SR Policy AFI. The Length field of the next-hop address specifies the next-hop address family. If the next-hop length is 4, then the next-hop is an IPv4 address. If the next-hop length is 16, then it is a global IPv6 address. If the next-hop length is 32, then it has a global IPv6 address followed by a link-local IPv6 address. The setting of the next-hop field and its attendant processing is governed by standard BGP procedures as

described in Section 3 of [RFC4760] and Section 3 of [RFC2545].

It is important to note that at any BGP speaker receiving BGP updates with SR Policy NLRIs, the SRPM processes only the best path as per the BGP best-path selection algorithm. In other words, this document leverages the existing BGP propagation and best-path selection rules. Details of the procedures are described in Section 4.

It has to be noted that if several CPs of the same SR Policy (Endpoint, Color) are signaled via BGP to a headend, then it is RECOMMENDED that each NLRI use a different distinguisher. If BGP has installed into the BGP table two advertisements whose respective NLRIs have the same Color and Endpoint, but different distinguishers, both advertisements are passed to the SRPM as different CPs along with their respective originator information (i.e., Autonomous System Number (ASN) and BGP Router-ID) as described in Section 2.4 of [RFC9256]. The ASN would be the ASN of the origin and the BGP Router-ID is determined in the following order:

- * From the Route Origin Community [RFC4360] if present and carrying an IP Address, or
- * As the BGP ORIGINATOR_ID [RFC4456] if present, or
- * As the BGP Router-ID of the peer from which the update was received as a last resort.

Section 2.9 of [RFC9256] specifies the selection of the active CP of the SR Policy by the SRPM based on the information provided to it by BGP.

2.2. SR Policy and Tunnel Encapsulation Attribute

The content of the SR Policy CP is encoded in the Tunnel Encapsulation Attribute defined in [RFC9012] using a Tunnel Type called the "SR Policy" type with code point 15. The use of the SR Policy Tunnel Type is applicable only for the AFI/SAFI pairs of (1/73, 2/73). This document specifies the use of the Tunnel Encapsulation Attribute with the SR Policy Tunnel Type and the use of any other Tunnel Type with the SR Policy SAFI MUST be considered malformed and handled by the "treat-as-withdraw" strategy [RFC7606].

The SR Policy Encoding structure is as follows:

SR Policy SAFI NLRI: <Distinguisher, Color, Endpoint>

Attributes:

```
Tunnel Encapsulation Attribute (23)
  Tunnel Type: SR Policy (15)
    Binding SID
    Preference
    Priority
    SR Policy Name
    SR Policy Candidate Path Name
    Explicit NULL Label Policy (ENLP)
    Segment List
      Weight
      Segment
      Segment
      ...
    ...
```

Figure 2: SR Policy Encoding

Where:

- * The SR Policy SAFI NLRI is defined in Section 2.1.

- * The Tunnel Encapsulation Attribute is defined in [RFC9012].
- * The Tunnel Type is set to 15.
- * Preference, Binding SID, Priority, SR Policy Name, SR Policy Candidate Path Name, ENLP, Segment-List, Weight, and Segment sub-TLVs are defined in Section 2.4.
- * Additional sub-TLVs may be defined in the future.

A Tunnel Encapsulation Attribute MUST NOT contain more than one TLV of type "SR Policy"; such updates MUST be considered malformed and handled by the "treat-as-withdraw" strategy [RFC7606].

BGP does not need to perform the validation of the tunnel (i.e., SR Policy) itself as indicated in Section 6 of [RFC9012]. The validation of the SR Policy information that is advertised using the sub-TLVs specified in Section 2.4 is performed by the SRPM.

2.3. Applicability of Tunnel Encapsulation Attribute Sub-TLVs

The Tunnel Egress Endpoint and Color sub-TLVs of the Tunnel Encapsulation Attribute, as defined in [RFC9012], are not utilized for SR Policy encodings. Consequently, their values are not relevant within the context of the SR Policy SAFI NLRI. If these sub-TLVs are present, a BGP speaker MUST ignore them and MAY remove them from the Tunnel Encapsulation Attribute during propagation.

Similarly, any other sub-TLVs, including those specified in [RFC9012], that do not have explicitly defined applicability to the SR Policy SAFI MUST be ignored by the BGP speaker and MAY be removed from the Tunnel Encapsulation Attribute during propagation.

2.4. SR Policy Sub-TLVs

This section specifies the sub-TLVs defined for encoding the information about the SR Policy Candidate Path.

Preference, Binding SID, SRv6 Binding SID, Segment-List, Priority, SR Policy Name, SR Policy Candidate Path Name, and Explicit NULL Label Policy are all optional sub-TLVs introduced for the BGP Tunnel Encapsulation Attribute [RFC9012] being defined in this section.

Weight and Segment are sub-TLVs of the Segment-List sub-TLV mentioned above.

An early draft version of this document included only the Binding SID sub-TLV that could be used for both SR-MPLS and SRv6 BSIDs. The SRv6 Binding SID TLV was introduced in later versions to support the advertisement of additional SRv6 capabilities without affecting backward compatibility for early implementations.

The fault management and error handling in the encoding of the sub-TLVs defined in this section are specified in Section 5. For the TLVs/sub-TLVs that are specified as single instance, only the first instance of that TLV/sub-TLV is used: the other instances MUST be ignored and MUST NOT be considered to be malformed.

None of the sub-TLVs defined in the following subsections have any effect on the BGP best-path selection or propagation procedures. These sub-TLVs are not used by the BGP path selection process and are instead passed on to SRPM as SR Policy Candidate Path information for further processing as described in Section 2 of [RFC9256].

The use of SR Policy sub-TLVs is applicable only for the AFI/SAFI

pairs of (1/73, 2/73). Future documents may extend their applicability to other AFI/SAFI.

2.4.1. Preference Sub-TLV

The Preference sub-TLV is used to carry the Preference of an SR Policy CP. The contents of this sub-TLV are used by the SRPM as described in Section 2.7 of [RFC9256].

The Preference sub-TLV is OPTIONAL; it MUST NOT appear more than once in the SR Policy encoding.

The Preference sub-TLV has the following format:




Figure 3: Preference Sub-TLV

Where:

Type: 12

Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 6.

Flags: 1 octet of flags. No flags are defined in this document. The Flags field MUST be set to zero on transmission and MUST be ignored on receipt.

RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

Preference: a 4-octet value indicating the Preference of the SR Policy CP as described in Section 2.7 of [RFC9256].


2.4.2. Binding SID Sub-TLV

The Binding SID sub-TLV is used to signal the BSID-related information of the SR Policy CP. The contents of this sub-TLV are used by the SRPM as described in Section 6 of [RFC9256].

The Binding SID sub-TLV is OPTIONAL; it MUST NOT appear more than once in the SR Policy encoding.

When the Binding SID sub-TLV is used to signal an SRv6 SID, the selection of the corresponding SRv6 Endpoint Behavior [RFC8986] to be instantiated is determined by the headend node. It is RECOMMENDED that the SRv6 Binding SID sub-TLV, as defined in Section 2.4.3, be used when signaling an SRv6 BSID for an SR Policy CP. The support for the use of this Binding SID sub-TLV for the signaling of an SRv6 BSID is retained primarily for backward compatibility with implementations that followed early draft versions of this document that had not defined the SRv6 Binding SID sub-TLV.

The Binding SID sub-TLV has the following format:




```

+-----+
|           Binding SID (variable, optional)           |
+-----+

```

Figure 4: Binding SID Sub-TLV

Where:

Type: 13

Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 18 when a SRv6 BSID is present, 6 when an SR-MPLS BSID is present, or 2 when no BSID is present.

Flags: 1 octet of flags. The following flags are defined in the registry "SR Policy Binding SID Flags" as described in Section 6.6:

```

  0 1 2 3 4 5 6 7
+-----+
|S|I|           |
+-----+

```

Figure 5: SR Policy Binding SID Flags

Where:

- * The S-Flag encodes the "Specified-BSID-Only" behavior. It is used by SRPM as described in Section 6.2.3 of [RFC9256].
- * The I-Flag encodes the "Drop-Upon-Invalid" behavior. It is used by SRPM as described in Section 8.2 of [RFC9256] to define a specific SR Policy forwarding behavior. The flag indicates that the SR Policy is to perform the "Drop-Upon-Invalid" behavior when no valid CP is available for this SR Policy. In this situation, the CP with the highest preference amongst those with the "Drop-Upon-Invalid" behavior is made active to drop traffic steered over the SR Policy.
- * The unassigned bits in the Flags field MUST be set to zero upon transmission and MUST be ignored upon receipt.

RESERVED: 1 octet of reserved bits. MUST be set to zero on transmission and MUST be ignored on receipt.

Binding SID: If the length is 2, then no BSID is present. If the length is 6, then the BSID is encoded in 4 octets using the format below. Traffic Class (TC), S, and TTL (Total of 12 bits) are RESERVED and MUST be set to zero and MUST be ignored.

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|           Label           | TC | S |           TTL           |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 6: Binding SID Label Encoding

The Label field is validated by the SRPM but MUST NOT contain the reserved MPLS label values (0-15). If the length is 18, then the BSID contains a 16-octet SRv6 SID.

2.4.3. SRv6 Binding SID Sub-TLV

The SRv6 Binding SID sub-TLV is used to signal the SRv6-BSID-related

information of an SR Policy CP. It enables the specification of the SRv6 Endpoint Behavior [RFC8986] to be instantiated on the headend node. The contents of this sub-TLV are used by the SRPM as described in Section 6 of [RFC9256].

The SRv6 Binding SID sub-TLV is OPTIONAL. More than one SRv6 Binding SID sub-TLV MAY be signaled in the same SR Policy encoding to indicate one or more SRv6 SIDs, each with potentially different SRv6 Endpoint Behaviors to be instantiated.

The SRv6 Binding SID sub-TLV has the following format:

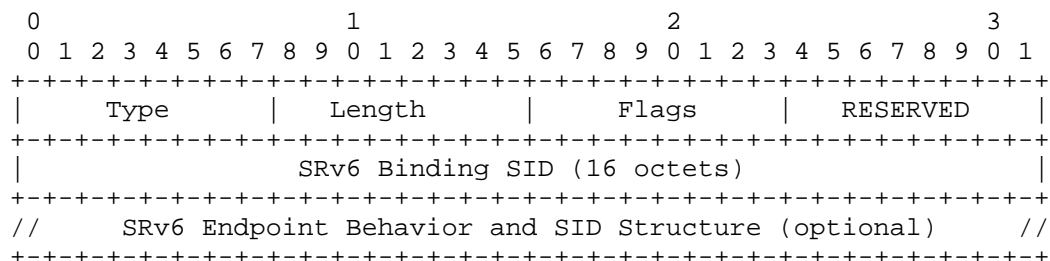


Figure 7: SRv6 Binding SID Sub-TLV

Where:

Type: 20

Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 26 when the SRv6 Endpoint Behavior and SID Structure is present; else, it MUST be 18.

Flags: 1 octet of flags. The following flags are defined in the registry "SR Policy SRv6 Binding SID Flags" as described in Section 6.7:




Figure 8: SR Policy SRv6 Binding SID Flags

Where:

- * The S-Flag encodes the "Specified-BSID-Only" behavior. It is used by SRPM as described in Section 6.2.3 of [RFC9256].
- * The I-Flag encodes the "Drop-Upon-Invalid" behavior. It is used by SRPM as described in Section 8.2 of [RFC9256].
- * The B-Flag, when set, indicates the presence of the "SRv6 Endpoint Behavior & SID Structure" encoding specified in Section 2.4.4.2.4.
- * The unassigned bits in the Flags field MUST be set to zero upon transmission and MUST be ignored upon receipt.

RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

SRv6 Binding SID: Contains a 16-octet SRv6 SID. The value 0 MAY be used when the controller wants to indicate the desired SRv6 Endpoint Behavior, SID Structure, or flags without specifying the BSID.

SRv6 Endpoint Behavior and SID Structure: Optional, as defined in Section 2.4.4.2.4. The SRv6 Endpoint Behavior and SID Structure MUST NOT be included when the SRv6 SID has not been included.

2.4.4. Segment List Sub-TLV

The Segment List sub-TLV encodes a single explicit path towards the Endpoint as described in Section 5.1 of [RFC9256]. The Segment List sub-TLV includes the elements of the paths (i.e., segments) as well as an optional Weight sub-TLV.

The Segment List sub-TLV may exceed 255 bytes in length due to a large number of segments. A 2-octet length is thus required. According to Section 2 of [RFC9012], the sub-TLV type defines the size of the Length field. Therefore, for the Segment List sub-TLV, a code point of 128 or higher is used.

The Segment List sub-TLV is OPTIONAL and MAY appear multiple times in the SR Policy encoding. The ordering of Segment List sub-TLVs does not matter since each sub-TLV encodes a Segment List.

The Segment List sub-TLV contains zero or more Segment sub-TLVs and MAY contain a Weight sub-TLV.

The Segment List sub-TLV has the following format:




Figure 9: Segment List Sub-TLV

Where:

Type: 128

Length: The total length (not including the Type and Length fields) of the sub-TLVs encoded within the Segment List sub-TLV in terms of octets.

RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

sub-TLVs currently defined:

- * An optional single Weight sub-TLV
- * Zero or more Segment sub-TLVs

Validation of an explicit path encoded by the Segment List sub-TLV is beyond the scope of BGP and performed by the SRPM as described in Section 5 of [RFC9256].

2.4.4.1. Weight Sub-TLV

The Weight sub-TLV specifies the weight associated with a given segment list. The contents of this sub-TLV are used only by the SRPM as described in Section 2.11 of [RFC9256].

The Weight sub-TLV is OPTIONAL; it MUST NOT appear more than once inside the Segment List sub-TLV.

The Weight sub-TLV has the following format:




Figure 10: Weight Sub-TLV

Where:

Type: 9

Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 6.

Flags: 1 octet of flags. No flags are defined in this document. The Flags field MUST be set to zero on transmission and MUST be ignored on receipt.

RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

Weight: 4 octets carrying an unsigned integer value indicating the weight associated with a segment list as described in Section 2.11 of [RFC9256]. A weight value of zero is invalid.

2.4.4.2. Segment Sub-TLVs

A Segment sub-TLV describes a single segment in a segment list (i.e., a single element of the explicit path). One or more Segment sub-TLVs constitute an explicit path of the SR Policy CP. The contents of these sub-TLVs are used only by the SRPM as described in Section 4 of [RFC9256].

The Segment sub-TLVs are OPTIONAL and MAY appear multiple times in the Segment List sub-TLV.

Section 4 of [RFC9256] defines several Segment Types:

Type A: SR-MPLS Label

Type B: SRv6 SID

Type C: IPv4 Prefix with optional SR Algorithm

Type D: IPv6 Global Prefix with optional SR Algorithm for SR-MPLS

Type E: IPv4 Prefix with Local Interface ID

Type F: IPv4 Addresses for link endpoints as Local, Remote pair

Type G: IPv6 Prefix and Interface ID for link endpoints as Local, Remote pair for SR-MPLS

Type H: IPv6 Addresses for link endpoints as Local, Remote pair for SR-MPLS

Type I: IPv6 Global Prefix with optional SR Algorithm for SRv6

Type J: IPv6 Prefix and Interface ID for link endpoints as Local, Remote pair for SRv6

Type K: IPv6 Addresses for link endpoints as Local, Remote pair for SRv6

The following subsections specify the sub-TLVs used for Segment Types A and B. The other segment types are specified in [RFC9831]. As specified in Section 5.1 of [RFC9256], a mix of SR-MPLS and SRv6 segments make the segment-list invalid.

2.4.4.2.1. Segment Type A

The Type A Segment sub-TLV encodes a single SR-MPLS SID. The format is as follows and is used to encode MPLS Label fields as specified in [RFC3032] and [RFC5462]:




Figure 11: Type A Segment Sub-TLV

Where:

Type: 1

Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 6.

Flags: 1 octet of flags as defined in Section 2.4.4.2.3.

RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

Label: 20 bits of label value.

TC: 3 bits of traffic class.

S: 1 bit of bottom-of-stack.

TTL: 1 octet of TTL.

The following applies to the Type-1 Segment sub-TLV:

- * The S bit MUST be zero upon transmission and MUST be ignored upon reception.
- * If the originator wants the receiver to choose the TC value, it sets the TC field to zero.
- * If the originator wants the receiver to choose the TTL value, it sets the TTL field to 255.
- * If the originator wants to recommend a value for these fields, it puts those values in the TC and/or TTL fields.
- * The receiver MAY override the originator's values for these fields. This would be determined by local policy at the receiver. One possible policy would be to override the fields only if the fields have the default values specified above.

2.4.4.2.2. Segment Type B

The Type B Segment sub-TLV encodes a single SRv6 SID. The format is

as follows:




Figure 12: Type B Segment Sub-TLV

Where:

Type: 13

Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 26 when the SRv6 Endpoint Behavior and SID Structure is present; else, it MUST be 18.

Flags: 1 octet of flags as defined in Section 2.4.4.2.3.

RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

SRv6 SID: 16 octets of IPv6 address.

SRv6 Endpoint Behavior and SID Structure: Optional, as defined in Section 2.4.4.2.4. The SRv6 Endpoint Behavior and SID Structure MUST NOT be included when the SRv6 SID has not been included.

The sub-TLV code point 2 defined for the advertisement of Segment Type B in the earlier draft versions of this document has been deprecated to avoid backward compatibility issues.

2.4.4.2.3. SR Policy Segment Flags

The Segment Type sub-TLVs described above may contain the following SR Policy Segment Flags in their Flags field. Also refer to Section 6.8:




Figure 13: SR Policy Segment Flags

Where:

- * When the V-Flag is set, it is used by SRPM for "SID verification" as described in Section 5.1 of [RFC9256].
- * When the B-Flag is set, it indicates the presence of the "SRv6 Endpoint Behavior & SID Structure" encoding specified in Section 2.4.4.2.4.
- * The unassigned bits in the Flags field MUST be set to zero upon transmission and MUST be ignored upon receipt.

The following applies to the Segment Flags:

- * V-Flag applies to all Segment Types.
- * B-Flag applies to Segment Type B. If B-Flag appears with Segment Type A, it MUST be ignored.

2.4.4.2.4. SRv6 Endpoint Behavior and SID Structure

The Segment Type sub-TLVs described above MAY contain the SRv6 Endpoint Behavior and SID Structure [RFC8986] encoding as described below:




Figure 14: SRv6 Endpoint Behavior and SID Structure

Where:

Endpoint Behavior: 2 octets. It carries the SRv6 Endpoint Behavior code point for this SRv6 SID as defined in Section 10.2 of [RFC8986]. When set with the value 0xFFFF (i.e., Opaque), the choice of SRv6 Endpoint Behavior is left to the headend.

Reserved: 2 octets of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

Locator Block Length: 1 octet. SRv6 SID Locator Block length in bits.

Locator Node Length: 1 octet. SRv6 SID Locator Node length in bits.

Function Length: 1 octet. SRv6 SID Function length in bits.

Argument Length: 1 octet. SRv6 SID Arguments length in bits.

The total of the locator block, locator node, function, and argument lengths MUST be less than or equal to 128.

2.4.5. Explicit NULL Label Policy Sub-TLV


To steer an unlabeled IP packet into an SR Policy for the MPLS data plane, it is necessary to push a label stack of one or more labels on that packet.

The Explicit NULL Label Policy (ENLP) sub-TLV is used to indicate whether an Explicit NULL Label [RFC3032] must be pushed on an unlabeled IP packet before any other labels.

If an ENLP sub-TLV is not present, the decision of whether to push an Explicit NULL label on a given packet is a matter of local configuration.

The ENLP sub-TLV is OPTIONAL; it MUST NOT appear more than once in the SR Policy encoding.

The contents of this sub-TLV are used by the SRPM as described in Section 4.1 of [RFC9256].



```

|      ENLP      |
+-----+

```

Figure 15: ENLP Sub-TLV

Where:

Type: 14

Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 3.

Flags: 1 octet of flags. No flags are defined in this document. The Flags field MUST be set to zero on transmission and MUST be ignored on receipt.

RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

ENLP (Explicit NULL Label Policy): Indicates whether Explicit NULL labels are to be pushed on unlabeled IP packets that are being steered into a given SR Policy. The following values have been currently defined for this field:

- 1: Push an IPv4 Explicit NULL label on an unlabeled IPv4 packet but do not push an IPv6 Explicit NULL label on an unlabeled IPv6 packet.
- 2: Push an IPv6 Explicit NULL label on an unlabeled IPv6 packet but do not push an IPv4 Explicit NULL label on an unlabeled IPv4 packet.
- 3: Push an IPv4 Explicit NULL label on an unlabeled IPv4 packet and push an IPv6 Explicit NULL label on an unlabeled IPv6 packet.
- 4: Do not push an Explicit NULL label.

This field can have one of the values as specified in Section 6.10. The ENLP unassigned values may be used for future extensions. Implementations adhering to this document MUST ignore the ENLP sub-TLV with unrecognized values (viz. other than 1 through 4). The behavior signaled in this sub-TLV MAY be overridden by local configuration by the network operator based on their deployment requirements. Section 4.1 of [RFC9256] describes the behavior on the headend for the handling of the Explicit NULL label.

2.4.6. SR Policy Priority Sub-TLV

An operator MAY set the SR Policy Priority sub-TLV to indicate the order in which the SR policies are recomputed upon topological change. The contents of this sub-TLV are used by the SRPM as described in Section 2.12 of [RFC9256].

The Priority sub-TLV is OPTIONAL; it MUST NOT appear more than once in the SR Policy encoding.

The Priority sub-TLV has the following format:

```

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 | Priority | RESERVED |
+-----+-----+-----+-----+

```


Figure 16: Priority Sub-TLV

Where:

Type: 15

Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value MUST be 2.

Priority: A 1-octet value indicating the priority as specified in Section 2.12 of [RFC9256].

RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

2.4.7. SR Policy Candidate Path Name Sub-TLV

An operator MAY set the SR Policy Candidate Path Name sub-TLV to attach a symbolic name to the SR Policy CP.

Usage of the SR Policy Candidate Path Name sub-TLV is described in Section 2.6 of [RFC9256].

The SR Policy Candidate Path Name sub-TLV may exceed 255 bytes in length due to a long name. A 2-octet length is thus required. According to Section 2 of [RFC9012], the sub-TLV type defines the size of the Length field. Therefore, for the SR Policy Candidate Path Name sub-TLV, a code point of 128 or higher is used.

It is RECOMMENDED that the size of the symbolic name for the CP be limited to 255 bytes. Implementations MAY choose to truncate long names to 255 bytes when signaling via BGP.

The SR Policy Candidate Path Name sub-TLV is OPTIONAL; it MUST NOT appear more than once in the SR Policy encoding.

The SR Policy Candidate Path Name sub-TLV has the following format:




Figure 17: SR Policy Candidate Path Name Sub-TLV

Where:

Type: 129

Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value is variable.

RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

SR Policy Candidate Path Name: Symbolic name for the SR Policy CP without a NULL terminator with encoding as specified in Section 2.6 of [RFC9256].

2.4.8. SR Policy Name Sub-TLV

An operator MAY set the SR Policy Name sub-TLV to associate a

symbolic name with the SR Policy for which the CP is being advertised via the SR Policy NLRI.

Usage of the SR Policy Name sub-TLV is described in Section 2.1 of [RFC9256].

The SR Policy Name sub-TLV may exceed 255 bytes in length due to a long SR Policy name. A 2-octet length is thus required. According to Section 2 of [RFC9012], the sub-TLV type defines the size of the Length field. Therefore, for the SR Policy Name sub-TLV, a code point of 128 or higher is used.

It is RECOMMENDED that the size of the symbolic name for the SR Policy be limited to 255 bytes. Implementations MAY choose to truncate long names to 255 bytes when signaling via BGP.

The SR Policy Name sub-TLV is OPTIONAL; it MUST NOT appear more than once in the SR Policy encoding.

The SR Policy Name sub-TLV has the following format:




Figure 18: SR Policy Name Sub-TLV

Where:

Type: 130

Length: Specifies the length of the value field (i.e., not including Type and Length fields) in terms of octets. The value is variable.

RESERVED: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.

SR Policy Name: Symbolic name for the SR Policy without a NULL terminator with encoding as specified in Section 2.1 of [RFC9256].

3. Color Extended Community

The Color Extended Community [RFC9012] is used to steer traffic corresponding to BGP routes into an SR Policy with matching Color value. The Color Extended Community MAY be carried in any BGP UPDATE message whose AFI/SAFI is 1/1 (IPv4 Unicast), 2/1 (IPv6 Unicast), 1/4 (IPv4 Labeled Unicast), 2/4 (IPv6 Labeled Unicast), 1/128 (VPN-IPv4 Labeled Unicast), 2/128 (VPN-IPv6 Labeled Unicast), or 25/70 (Ethernet VPN, usually known as EVPN). Use of the Color Extended Community in BGP UPDATE messages of other AFI/SAFIs is not covered by [RFC9012]; hence, it is outside the scope of this document as well.

Two bits from the Flags field of the Color Extended Community are used as follows to support the requirements of Color-Only steering as specified in Section 8.8 of [RFC9256]:




Figure 19: Color Extended Community Flags

The C and O bits together form the Color-Only Type field, which indicates the various matching criteria between the BGP Next Hop (NH) and the SR Policy Endpoint in addition to the matching of the Color value. The following types are defined:

Type 0 (bits 00): Specific Endpoint Match. Request a match for the Endpoint that is the BGP NH.

Type 1 (bits 01): Specific or Null Endpoint Match. Request a match for either the Endpoint that is the BGP NH or a null Endpoint (e.g., a default gateway).

Type 2 (bits 10): Specific, Null, or Any Endpoint Match. Request a match for either the Endpoint that is the BGP NH or a null or any Endpoint.

Type 3 (bits 11): Reserved for future use and SHOULD NOT be used. Upon reception, an implementation MUST treat it like Type 0.

The details of the SR Policy steering mechanisms based on these Color-Only types are specified in Section 8.8 of [RFC9256].

One or more Color Extended Communities MAY be associated with a BGP route update. Sections 8.4.1, 8.5.1, and 8.8.2 of [RFC9256] specify the steering behaviors over SR Policies when multiple Color Extended Communities are associated with a BGP route.

4. SR Policy Operations

As mentioned in Section 1, BGP is not the actual consumer of an SR Policy NLRI. BGP is in charge of the origination and propagation of the SR Policy NLRI, but its installation and use are outside the scope of BGP. The details of SR Policy installation and use are specified in [RFC9256].

4.1. Advertisement of SR Policies

Typically, but not limited to, an SR Policy is computed by a controller or a Path Computation Engine (PCE) and originated by a BGP speaker on its behalf.

Multiple SR Policy NLRIs may be present with the same <Color, Endpoint> tuple but with different distinguishers when these SR policies are intended for different headends.

The distinguisher of each SR Policy NLRI prevents undesired BGP route selection among these SR Policy NLRIs and allows their propagation across RRs [RFC4456].

Moreover, one or more route targets SHOULD be attached to the advertisement, where each route target identifies one or more intended headends for the advertised SR Policy update.

If no route target is attached to the SR Policy NLRI, then it is assumed that the originator sends the SR Policy update directly (e.g., through a BGP session) to the intended receiver. In such a case, the NO_ADVERTISE community [RFC1997] MUST be attached to the SR Policy update (see further details in Section 4.2.3).

4.2. Reception of an SR Policy NLRI

On reception of an SR Policy NLRI, a BGP speaker first determines if it is valid as described in Section 4.2.1; then, the BGP speaker

performs the decision process for selection of the best route (Section 9.1 of [RFC4271]). The key difference from the base BGP decision process is that BGP does not download the selected best routes of the SR Policy SAFI into the forwarding; instead, it considers them "usable" for passing on to the SRPM for further processing as described in Section 4.2.2. The selected best route is "propagated" (Section 9.1.3 of [RFC4271]) as described in Section 4.2.3, irrespective of its "usability" by the local router.

4.2.1. Validation of an SR Policy NLRI

When a BGP speaker receives an SR Policy NLRI from a neighbor, it MUST first perform validation based on the following rules in addition to the validation described in Section 5:

- * The SR Policy NLRI MUST include a distinguisher, Color, and Endpoint field that implies that the length of the NLRI MUST be either 12 or 24 octets (depending on the address family of the Endpoint).
- * The SR Policy update MUST have either the NO_ADVERTISE community, at least one Route Target extended community in IPv4-address format, or both. If a router supporting this specification receives an SR Policy update with no Route Target extended communities and no NO_ADVERTISE community, the update MUST be considered to be malformed.
- * The Tunnel Encapsulation Attribute MUST be attached to the BGP UPDATE message and MUST have a Tunnel Type TLV set to SR Policy (code point is 15).

A router that receives an SR Policy update that is not valid according to these criteria MUST treat the update as malformed, and the SR Policy CP MUST NOT be passed to the SRPM.

4.2.2. Eligibility for Local Use of an SR Policy NLRI

An SR Policy NLRI update that does not have a Route Target extended community but does have the NO_ADVERTISE community is considered usable.

If one or more route targets are present, then at least one route target MUST match the BGP Identifier of the receiver for the update to be considered usable. The BGP Identifier is defined in [RFC4271] as a 4-octet IPv4 address and is updated by [RFC6286] as a 4-octet, unsigned, non-zero integer. Therefore, the Route Target extended community MUST be of the same format.

If one or more route targets are present, and none matches the local BGP Identifier, then, while the SR Policy NLRI is valid, the SR Policy NLRI is not usable on the receiver node.

When the SR Policy tunnel type includes any sub-TLV that is unrecognized or unsupported, the update SHOULD NOT be considered usable. An implementation MAY provide an option for ignoring unsupported sub-TLVs.

Once BGP on the receiving node has determined that the SR Policy NLRI is usable, it passes the SR Policy CP to the SRPM. Note that, along with the CP details, BGP also passes the originator information for breaking ties in the CP selection process as described in Section 2.4 of [RFC9256].

When an update for an SR Policy NLRI results in its becoming unusable, BGP MUST delete its corresponding SR Policy CP from the SRPM.

The SRPM applies the rules defined in Section 2 of [RFC9256] to determine whether the SR Policy CP is valid and to select the active CP for a given SR Policy.

4.2.3. Propagation of an SR Policy

SR Policy NLRIs that have the NO_ADVERTISE community attached to them MUST NOT be propagated.

By default, a BGP node receiving an SR Policy NLRI MUST NOT propagate it to any External BGP (EBGP) neighbor. An implementation MAY provide an explicit configuration to override this and enable the propagation of valid SR Policy NLRIs to specific EBGP neighbors where the SR domain comprises multiple ASes within a single service provider domain (see Section 7 for details).

A BGP node advertises a received SR Policy NLRI to its Internal BGP (IBGP) neighbors according to normal IBGP propagation rules.

By default, a BGP node receiving an SR Policy NLRI SHOULD NOT remove the Route Target extended community before propagation. An implementation MAY provide support for configuration to filter and/or remove the Route Target extended community before propagation.

A BGP node MUST NOT alter the SR Policy information carried in the Tunnel Encapsulation Attribute during propagation.

5. Error Handling and Fault Management

This section describes the error-handling actions, as described in [RFC7606], that are to be performed for the handling of the BGP UPDATE messages for the BGP SR Policy SAFI.

A BGP speaker MUST perform the following syntactic validation of the SR Policy NLRI to determine if it is malformed. This includes the validation of the length of each NLRI and the total length of the MP_REACH_NLRI and MP_UNREACH_NLRI attributes. It also includes the validation of the consistency of the NLRI length with the AFI and the endpoint address as specified in Section 2.1.

When the error determined allows for the router to skip the malformed NLRI(s) and continue the processing of the rest of the BGP UPDATE message, then it MUST handle such malformed NLRIs as 'treat-as-withdraw'. In other cases, where the error in the NLRI encoding results in the inability to process the BGP UPDATE message (e.g., length-related encoding errors), then the router SHOULD handle such malformed NLRIs as "AFI/SAFI disable" when other AFI/SAFIs besides SR Policy are being advertised over the same session. Alternately, the router MUST perform "session reset" when the session is only being used for SR Policy or when a "AFI/SAFI disable" action is not possible.

The validation of the TLVs/sub-TLVs introduced in this document and defined in their respective subsections of Section 2.4 MUST be performed to determine if they are malformed or invalid. The validation of the Tunnel Encapsulation Attribute itself and the other TLVs/sub-TLVs specified in Section 13 of [RFC9012] MUST be done as described in that document. In case of any error detected, either at the attribute or its TLV/sub-TLV level, the "treat-as-withdraw" strategy MUST be applied. This is because an SR Policy update without a valid Tunnel Encapsulation Attribute (comprised of all valid TLVs/sub-TLVs) is not usable.

An SR Policy update that is determined not to be valid (and, therefore, malformed) based on the rules described in Section 4.2.1

MUST be handled by the "treat-as-withdraw" strategy.

The validation of the individual fields of the TLVs/sub-TLVs defined in Section 2.4 are beyond the scope of BGP as they are handled by the SRPM as described in the individual TLV/sub-TLV subsections. A BGP implementation MUST NOT perform semantic verification of such fields nor consider the SR Policy update to be invalid or not usable based on such validation.

An implementation SHOULD log any errors found during the above validation for further analysis.

6. IANA Considerations

This document uses code point allocations from the following existing registries in the "Subsequent Address Family Identifiers (SAFI) Parameters" registry group:

- * The "SAFI Values" registry

This document uses code point allocations from the following existing registries in the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group:

- * The "BGP Tunnel Encapsulation Attribute Tunnel Types" registry
- * The "BGP Tunnel Encapsulation Attribute Sub-TLVs" registry
- * The "Color Extended Community Flags" registry

This document creates the following new registries in the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group:

- * The "SR Policy Segment List Sub-TLVs" registry
- * The "SR Policy Binding SID Flags" registry
- * The "SR Policy SRv6 Binding SID Flags" registry
- * The "SR Policy Segment Flags" registry
- * The "Color Extended Community Color-Only Types" registry

This document creates the following new registry in the "Segment Routing" registry group:

- * The "SR Policy ENLP Values" registry

6.1. Subsequent Address Family Identifiers (SAFI) Parameters

This document registers a SAFI code point in the "SAFI Values" registry of the "Subsequent Address Family Identifiers (SAFI) Parameters" registry group as follows:

Value	Description	Reference
73	SR Policy SAFI	RFC 9830

Table 1: BGP SAFI Code Point

6.2. BGP Tunnel Encapsulation Attribute Tunnel Types

This document registers a Tunnel Type code point in the "BGP Tunnel Encapsulation Attribute Tunnel Types" registry under the "Border

Gateway Protocol (BGP) Tunnel Encapsulation" registry group.

Value	Description	Reference
15	SR Policy	RFC 9830

Table 2: Tunnel Type Code Point

6.3. BGP Tunnel Encapsulation Attribute Sub-TLVs

This document defines sub-TLVs in the "BGP Tunnel Encapsulation Attribute Sub-TLVs" registry under the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group.

Value	Description	Reference	Change Controller
12	Preference sub-TLV	RFC 9830	IETF
13	Binding SID sub-TLV	RFC 9830	IETF
14	ENLP sub-TLV	RFC 9830	IETF
15	Priority sub-TLV	RFC 9830	IETF
20	SRv6 Binding SID sub-TLV	RFC 9830	IETF
128	Segment List sub-TLV	RFC 9830	IETF
129	SR Policy Candidate Path Name sub-TLV	RFC 9830	IETF
130	SR Policy Name sub-TLV	RFC 9830	IETF

Table 3: BGP Tunnel Encapsulation Attribute Sub-TLV Code Points

6.4. Color Extended Community Flags

This document defines the use of 2 bits in the "Color Extended Community Flags" registry under the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group.

Bit Position	Description	Reference
0-1	Color-only Types Field	RFC 9830

Table 4: Color Extended Community Flag Bits

6.5. SR Policy Segment List Sub-TLVs

This document creates a new registry called "SR Policy Segment List Sub-TLVs" under the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group. The registration policy of this registry is "IETF Review" (see [RFC8126]).

The following initial sub-TLV code points are assigned by this document:

Value	Description	Reference
-------	-------------	-----------

0	Reserved	RFC 9830	
+-----+	+-----+	+-----+	+-----+
1	Type A Segment sub-TLV	RFC 9830	
+-----+	+-----+	+-----+	+-----+
2	Deprecated	RFC 9830	
+-----+	+-----+	+-----+	+-----+
3-8	Unassigned		
+-----+	+-----+	+-----+	+-----+
9	Weight sub-TLV	RFC 9830	
+-----+	+-----+	+-----+	+-----+
10	Deprecated	RFC 9830	
+-----+	+-----+	+-----+	+-----+
11	Deprecated	RFC 9830	
+-----+	+-----+	+-----+	+-----+
12	Deprecated	RFC 9830	
+-----+	+-----+	+-----+	+-----+
13	Type B Segment sub-TLV	RFC 9830	
+-----+	+-----+	+-----+	+-----+
14-255	Unassigned		
+-----+	+-----+	+-----+	+-----+

Table 5: SR Policy Segment List Sub-TLV
Code Points

6.6. SR Policy Binding SID Flags

This document creates a new registry called "SR Policy Binding SID Flags" under the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group. The registration policy of this registry is "Standards Action" (see [RFC8126]).

The following flags are defined:

+=====+	+=====+	+=====+	+=====+
Bit	Description	Reference	
+=====+	+=====+	+=====+	+=====+
0	Specified-BSID-Only Flag (S-Flag)	RFC 9830	
+-----+	+-----+	+-----+	+-----+
1	Drop-Upon-Invalid Flag (I-Flag)	RFC 9830	
+-----+	+-----+	+-----+	+-----+
2-7	Unassigned		
+-----+	+-----+	+-----+	+-----+

Table 6: SR Policy Binding SID Flags

6.7. SR Policy SRv6 Binding SID Flags

This document creates a new registry called "SR Policy SRv6 Binding SID Flags" under the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group. The registration policy of this registry is "Standards Action" (see [RFC8126]).

The following flags are defined:

+=====+	+=====+	+=====+	+=====+
Bit	Description	Reference	
+=====+	+=====+	+=====+	+=====+
0	Specified-BSID-Only Flag (S-Flag)	RFC 9830	
+-----+	+-----+	+-----+	+-----+
1	Drop-Upon-Invalid Flag (I-Flag)	RFC 9830	
+-----+	+-----+	+-----+	+-----+
2	SRv6 Endpoint Behavior & SID Structure Flag (B-Flag)	RFC 9830	
+-----+	+-----+	+-----+	+-----+
3-7	Unassigned		
+-----+	+-----+	+-----+	+-----+

Table 7: SR Policy SRv6 Binding SID Flags

6.8. SR Policy Segment Flags

This document creates a new registry called "SR Policy Segment Flags" under the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group. The registration policy of this registry is "IETF Review" (see [RFC8126]).

The following flags are defined:

Bit	Description	Reference
0	Segment Verification Flag (V-Flag)	RFC 9830
1-2	Unassigned	
3	SRv6 Endpoint Behavior & SID Structure Flag (B-Flag)	RFC 9830
4-7	Unassigned	

Table 8: SR Policy Segment Flags

6.9. Color Extended Community Color-Only Types

This document creates a new registry called "Color Extended Community Color-Only Types" under the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group for assignment of code points (values 0 through 3) in the Color-Only Type field of the Color Extended Community Flags field. The registration policy of this registry is "Standards Action" (see [RFC8126]).

The following types are defined:

Type	Description	Reference
0	Specific Endpoint Match	RFC 9830
1	Specific or Null Endpoint Match	RFC 9830
2	Specific, Null, or Any Endpoint Match	RFC 9830
3	Unassigned	RFC 9830

Table 9: Color Extended Community Color-Only Types

6.10. SR Policy ENLP Values

IANA will maintain a new registry under the "Segment Routing" registry group with the registration policy of "Standards Action" (see [RFC8126]). The new registry is called "SR Policy ENLP Values" and contains the code points allocated to the ENLP field defined in Section 2.4.5. The registry contains the following code points:

Code Point	Description	Reference
0	Reserved	RFC 9830

1	Push an IPv4 Explicit NULL label on an unlabeled IPv4 packet but do not push an IPv6 Explicit NULL label on an unlabeled IPv6 packet	RFC 9830
2	Push an IPv6 Explicit NULL label on an unlabeled IPv6 packet but do not push an IPv4 Explicit NULL label on an unlabeled IPv4 packet	RFC 9830
3	Push an IPv6 Explicit NULL label on an unlabeled IPv6 packet and push an IPv4 Explicit NULL label on an unlabeled IPv4 packet	RFC 9830
4	Do not push an Explicit NULL label	RFC 9830
5-255	Unassigned	

Table 10: SR Policy ENLP Values

7. Security Considerations

The security mechanisms of the base BGP security model apply to the extensions described in this document as well. See the Security Considerations section of [RFC4271] for a discussion of BGP security. Also, refer to [RFC4272] and [RFC6952] for analysis of security issues for BGP.

The BGP SR Policy extensions specified in this document enable traffic engineering and service programming use cases within an SR domain as described in [RFC9256]. SR operates within a trusted SR domain [RFC8402]; its security considerations also apply to BGP sessions when carrying SR Policy information. The SR Policies distributed by BGP are expected to be used entirely within this trusted SR domain, which comprises a single AS or multiple ASes / domains within a single provider network. Therefore, precaution is necessary to ensure that the SR Policy information advertised via BGP sessions is limited to nodes in a secure manner within this trusted SR domain. BGP peering sessions for address families other than those that use the SR Policy SAFI may be set up to routers outside the SR domain. The isolation of BGP SR Policy SAFI peering sessions may be used to ensure that the SR Policy information is not advertised by accident or in error to an EBGp peering session outside the SR domain.

Additionally, it may be a consideration that the export of SR Policy information, as described in this document, constitutes a risk to confidentiality of mission-critical or commercially sensitive information about the network (more specifically endpoint/node addresses, SR SIDs, and the SR Policies deployed). BGP peerings are not automatic and require configuration; thus, it is the responsibility of the network operator to ensure that only trusted nodes (that include both routers and controller applications) within the SR domain are configured to receive such information.

8. Manageability Considerations

The specification of BGP models is an ongoing work based on [BGP-YANG-MODEL]; its future extensions are expected to cover the SR Policy SAFI. Existing BGP operational procedures also apply to the SAFI specified in this document. The management, operations, and monitoring of BGP speakers and the SR Policy SAFI sessions between them are not very different from other BGP sessions and can be

managed using the same data models.

The YANG data model for the operation and management of SR Policies [SR-POLICY-YANG] reports the SR Policies provisioned via BGP SR Policy SAFI along with their operational states.

9. References

9.1. Normative References

- [RFC1997] Chandra, R., Traina, P., and T. Li, "BGP Communities Attribute", RFC 1997, DOI 10.17487/RFC1997, August 1996, <<https://www.rfc-editor.org/info/rfc1997>>.
- [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>>.
- [RFC2545] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing", RFC 2545, DOI 10.17487/RFC2545, March 1999, <<https://www.rfc-editor.org/info/rfc2545>>.
- [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>>.
- [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>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.
- [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>>.
- [RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, DOI 10.17487/RFC5462, February 2009, <<https://www.rfc-editor.org/info/rfc5462>>.
- [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>>.
- [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>>.
- [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>>.

- [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>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

9.2. Informative References

- [BGP-LS-SR-POLICY] Previdi, S., Talaulikar, K., Ed., Dong, J., Gredler, H., and J. Tantsura, "Advertisement of Segment Routing Policies using BGP Link-State", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-sr-policy-17, 6 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-17>>.
- [BGP-YANG-MODEL] Jethanandani, M., Patel, K., Hares, S., and J. Haas, "YANG Model for Border Gateway Protocol (BGP-4)", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-model-18, 21 October 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-model-18>>.
- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <<https://www.rfc-editor.org/info/rfc4272>>.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <<https://www.rfc-editor.org/info/rfc4456>>.
- [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, <<https://www.rfc-editor.org/info/rfc6952>>.
- [RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552,

DOI 10.17487/RFC9552, December 2023,
<<https://www.rfc-editor.org/info/rfc9552>>.

[RFC9831] Talaulikar, K., Ed., Filsfils, C., Previdi, S., Mattes, P., and D. Jain, "Segment Type Extensions for BGP Segment Routing (SR) Policy", RFC 9831, DOI 10.17487/RFC9831, September 2025, <<https://www.rfc-editor.org/info/rfc9831>>.

[SR-POLICY-YANG] Raza, K., Ed., Saleh, T., Shunwan, Z., Voyer, D., Durrani, M., Matsushima, S., and V. Beeram, "YANG Data Model for Segment Routing Policy", Work in Progress, Internet-Draft, draft-ietf-spring-sr-policy-yang-05, 25 May 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-policy-yang-05>>.

Acknowledgments

The authors of this document would like to thank Shyam Sethuram, John Scudder, Przemyslaw Krol, Alex id Bogdanov, Nandan Saha, Bruno Decraene, Gurusiddesh Nidasesi, Kausik Majumdar, Zafar Ali, Swadesh Agarwal, Jakob Heitz, Viral Patel, Peng Shaofu, Cheng Li, Martin Vigoureux, John Scudder, Vincent Roca, Brian Haberman, Mohamed Boucadair, Shunwan Zhuang, Andrew Alston, Jeffrey (Zhaohui) Zhang, Nagendra Nainar, Rajesh Melarcode Venkateswaran, Nat Kao, Boris Hassanov, Vincent Roca, Russ Housley, and Dan Romascanu for their comments and review of this document. The authors would like to thank Susan Hares for her detailed shepherd review that helped in improving the document.

Contributors

Eric Rosen
Juniper Networks
United States of America
Email: erosen@juniper.net

Arjun Sreekantiah
Cisco Systems
United States of America
Email: asreekan@cisco.com

Acee Lindem
Cisco Systems
United States of America
Email: acee@cisco.com

Siva Sivabalan
Cisco Systems
United States of America
Email: msiva@cisco.com

Imtiyaz Mohammad
Arista Networks
India
Email: imtiyaz@arista.com

Gaurav Dawra
Cisco Systems
United States of America
Email: gdawra.ietf@gmail.com

Peng Shaofu
ZTE Corporation
China
Email: peng.shaofu@zte.com.cn

Steven Lin
Calix
United States of America
Email: steven.lin@calix.com

Authors' Addresses

Stefano Previdi
Huawei Technologies
Italy
Email: stefano@previdi.net

Clarence Filsfils
Cisco Systems
Brussels
Belgium
Email: cfilsfil@cisco.com

Ketan Talaulikar (editor)
Cisco Systems
India
Email: ketant.ietf@gmail.com

Paul Mattes
Microsoft
One Microsoft Way
Redmond, WA 98052
United States of America
Email: pamattes@microsoft.com

Dhanendra Jain
Google
Email: dhanendra.ietf@gmail.com