Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: 7 September 2025

S. Previdi
Individual
K. Talaulikar, Ed.
Cisco Systems
J. Dong
Huawei Technologies
H. Gredler
RtBrick Inc.
J. Tantsura
Nvidia

6 March 2025

# Advertisement of Segment Routing Policies using BGP Link-State draft-ietf-idr-bgp-ls-sr-policy-17

## Abstract

This document describes a mechanism to collect the Segment Routing Policy information that is locally available in a node and advertise it into BGP Link-State (BGP-LS) updates. Such information can be used by external components for path computation, re-optimization, service placement, network visualization, etc.

## Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 7 September 2025.

## Copyright Notice

Table of Contents

## 1. Introduction

SR Policy architecture details are specified in [RFC9256]. An SR Policy comprises one or more candidate paths of which at a given time one and only one may be active (i.e., installed in forwarding and usable for steering of traffic). Each candidate path in turn may have one or more SID-List of which one or more SID-List may be active. When multiple SID-Lists are active then traffic is load balanced over them. This document covers the advertisement of state information at the individual SR Policy candidate path level.

SR Policies are generally instantiated at the head-end and are based on either local configuration or controller-based programming of the node using various APIs and protocols (e.g., PCEP or BGP).

In many network environments, the configuration, and state of each SR Policy that is available in the network is required by controllers. Such controllers, that are aware of both topology and SR Policy state information, allow the network operator to optimize several functions and operations in their networks.

One example of a controller is the stateful Path Computation Element (PCE) [RFC8231], which could provide benefits in path optimization. While some extensions are proposed in the Path Computation Element Communication Protocol (PCEP) for the Path Computation Clients (PCCs) to report the LSP states to the PCE, this mechanism may not be applicable in a management-based PCE architecture as specified in section 5.5 of [RFC4655]. As illustrated in the figure below, the PCC is not an LSR in the routing domain, thus the head-end nodes of the SR Policies may not implement the PCEP protocol. In this case, a general mechanism to collect the SR Policy states from the ingress LERs is needed. This document proposes an SR Policy state collection mechanism complementary to the mechanism defined in [RFC8231].

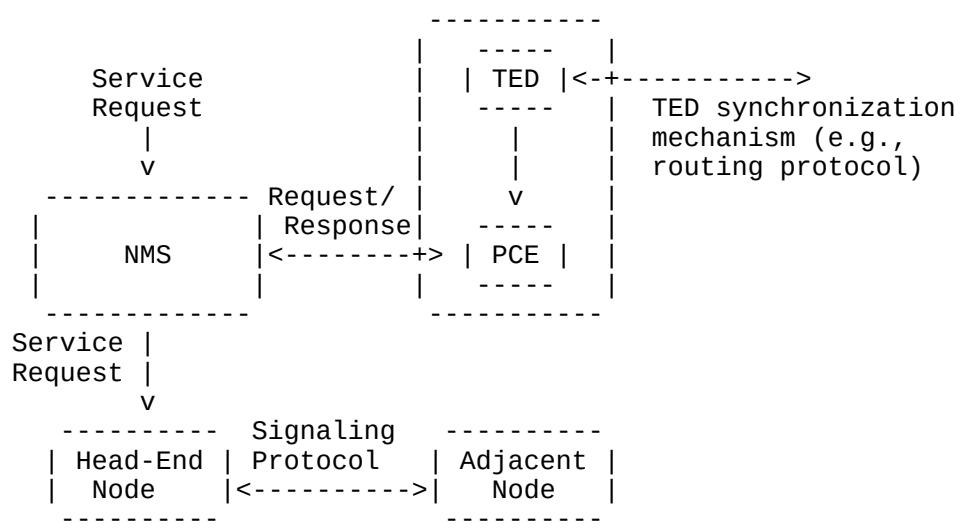


Figure 1  Management-Based PCE Usage

In networks with composite PCE nodes as specified in section 5.1 of [RFC4655], PCE is implemented on several routers in the network, and the PCCs in the network can use the mechanism described in [RFC8231] to report the SR Policy information to the PCE nodes.  An external component may also need to collect the SR Policy information from all the PCEs in the network to obtain a global view of the state of all SR Policy paths in the network.

In multi-area or multi-AS scenarios, each area or AS can have a child PCE to collect the SR Policies in its domain, in addition, a parent PCE needs to collect SR Policy information from multiple child PCEs to obtain a global view of SR Policy paths inside and across the domains involved.

In another network scenario, a centralized controller is used for service placement.  Obtaining the SR Policy state information is quite important for making appropriate service placement decisions with the purpose of both meeting the application’s requirements and utilizing network resources efficiently.

The Network Management System (NMS) may need to provide global visibility of the SR Policies in the network as part of the network visualization function.

BGP has been extended to distribute link-state and traffic engineering information to external components [RFC9552].  Using the same protocol to collect SR Policy and state information is desirable

for these external components since this avoids introducing multiple protocols for network topology information collection. This document describes a mechanism to distribute SR Policy information (both SR-MPLS, and SRv6 [RFC8402]) to external components using BGP-LS and covers both explicit and dynamic candidate paths. The advertisements of composite candidate path is outside the scope of this document.

The BGP-LS Producer [RFC9552] that is originating the advertisement of SR Policy information can be either:

* a SR Policy headend node, or

* a PCE which is receiving the SR Policy information from its PCCs (i.e., SR Policy headend nodes) via PCEP

The extensions specified in this document complement the BGP SR Policy SAFI [I-D.ietf-idr-sr-policy-safi] and [I-D.ietf-idr-bgp-sr-segtypes-ext] that are used to advertise SR Policies from controllers to the headend routers using BGP by enabling the reporting of the operational state of those SR Policies back from the headend to the controllers.

While this document focuses on SR Policies, [I-D.ietf-idr-bgp-ls-te-path] introduces further extension to support other TE Paths such as MPLS-TE LSPs.

The encodings specified in this document (specifically in Section 4 and Section 5) make use of flags that convey various types of information of the SR Policy. The document uses the term "set" to indicate that the value of a flag bit is 1 and the term "clear" when the value is 0.

### 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. Carrying SR Policy Information in BGP

The "Link-State NLRI" defined in [RFC9552] is extended to carry the SR Policy information. New TLVs carried in the BGP Link-State Attribute defined in [RFC9552] are also defined to carry the attributes of an SR Policy in the subsequent sections.

The format of "Link-State NLRI" is defined in [RFC9552] as follows:

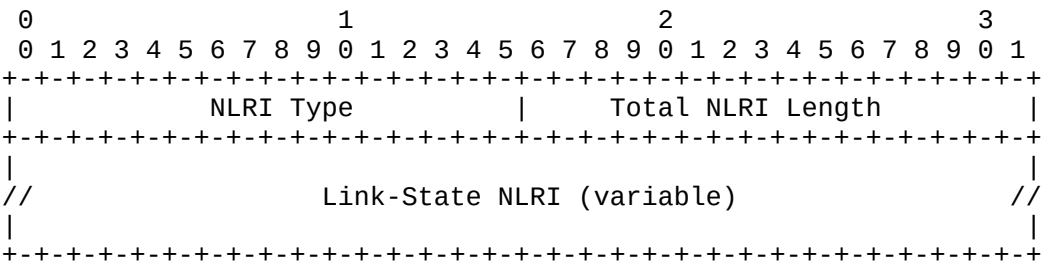


Figure 2  BGP-LS NLRI Format

An additional "NLRI Type" known as SR Policy Candidate Path NLRI (value 5) is defined for the advertisement of SR Policy Information.

This SR Policy Candidate Path NLRI is used to report the state details of individual SR Policy Candidate paths along with their underlying segment lists.

## 3.  SR Policy Candidate Path NLRI Type

This document defines SR Policy Candidate Path NLRI Type with its format as shown in the following figure:

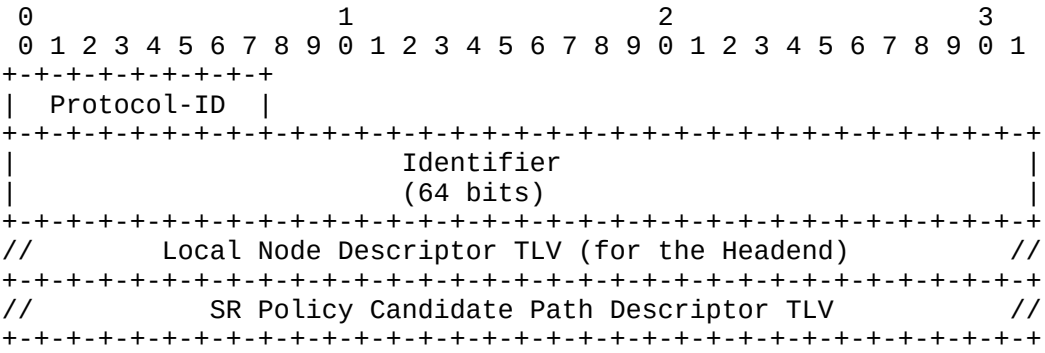


Figure 3  SR Policy Candidate Path NLRI Format

Where:

* Protocol-ID field specifies the component that owns the SR Policy state in the advertising node.  An additional Protocol-ID "Segment Routing" (value 9) is introduced by this document that MUST be used for advertisement of SR Policies.

* "Identifier" is an 8 octet value as defined in section 5.2 of [RFC9552].

* "Local Node Descriptor" (TLV 256) [RFC9552] is used as specified further in this section.

* The SR Policy Candidate Path Descriptor TLV is specified in Section 4.

The Local Node Descriptor TLV carries information that only identifies the headend node of the SR Policy irrespective of whether the BGP-LS Producer is a headend or a PCE node.

The Local Node Descriptor TLV MUST include at least one of the following Node Descriptor TLVs:

* IPv4 Router-ID of Local Node (TLV 1028) [RFC9552], which identifies the headend node of the SR Policy as specified in section 2.1 of [RFC9256].

* IPv6 Router-ID of Local Node (TLV 1029) [RFC9552], which identifies the headend node of the SR Policy as specified in section 2.1 of [RFC9256].

The following sub-sections describe the encoding of sub-TLVs within the Local Node Descriptor TLV depending on which node is the BGP-LS Producer.

## 3.1. SR Policy Headend as BGP-LS Producer

The Local Node Descriptor TLV MUST include the following Node Descriptor TLVs when the headend node is the BGP-LS Producer:

* BGP Router-ID (TLV 516) [RFC9086], which contains a valid BGP Identifier of the headend node of the SR Policy.

* Autonomous System Number (TLV 512) [RFC9552], which contains the ASN (or AS Confederation Identifier (ASN) [RFC5065], if confederations are used) of the headend node of the SR Policy.

The Local Node Descriptor TLV MAY include the following Node Descriptor TLVs when the headend node is the BGP-LS Producer:

* BGP Confederation Member (TLV 517) [RFC9086], which contains the ASN of the confederation member (i.e. Member-AS Number), if BGP confederations are used, the headend node of the SR Policy.

* Other Node Descriptors as defined in [RFC9552] to identify the headend node of the SR Policy. The determination of whether the IGP Router-ID sub-TLV (TLV 515) contains a 4-octet OSPF Router-ID or a 6-octet ISO System-ID is to be done based on the length of that sub-TLV since the Protocol-ID in the NLRI is always going to be "Segment Routing".

### 3.2. PCE as BGP-LS Producer

The PCE node MUST NOT include its identifiers in the Node Descriptor TLV in the NLRI as the Node Descriptor TLV MUST only carry the identifiers of the SR Policy headend.

The Local Node Descriptor TLV MAY include the following Node Descriptor TLVs when the PCE node is the BGP-LS Producer and it has this information about the headend (e.g., as part of its topology database):

* BGP Router-ID (TLV 516) [RFC9086], which contains a valid BGP Identifier of the headend node of the SR Policy.

* Autonomous System Number (TLV 512) [RFC9552], which contains the ASN (or AS Confederation Identifier (ASN) [RFC5065], if confederations are used) of the headend node of the SR Policy.

* BGP Confederation Member (TLV 517) [RFC9086], which contains the ASN of the confederation member (i.e. Member-AS Number), if BGP confederations are used, the headend node of the SR Policy.

* Other Node Descriptors as defined in [RFC9552] to identify the headend node of the SR Policy. The determination of whether the IGP Router-ID sub-TLV (TLV 515) contains a 4-octet OSPF Router-ID or a 6-octet ISO System-ID is to be done based on the length of that sub-TLV since the Protocol-ID in the NLRI is always going to be "Segment Routing".

When a Path Computation Element (PCE) node is functioning as the BGP-LS Producer on behalf of one or more headends, it MAY include its own BGP Router-ID (TLV 516), Autonomous System Number (TLV 512), or BGP Confederation Member (TLV 517) in the BGP-LS Attribute.

## 4. SR Policy Candidate Path Descriptor

The SR Policy Candidate Path Descriptor TLV identifies a Segment Routing Policy candidate path as defined in [RFC9256]. It is a mandatory TLV for SR Policy Candidate Path NLRI type. The TLV has the following format:

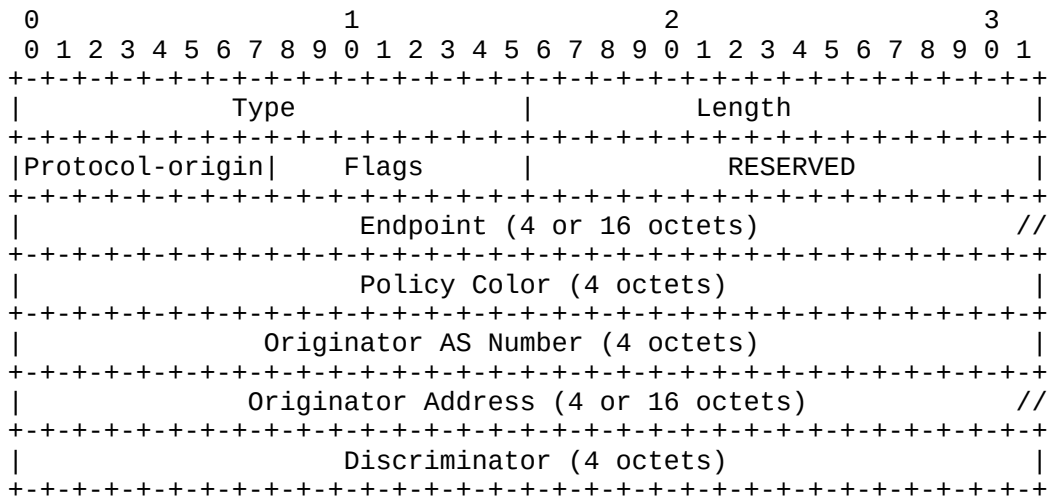


Figure 4  SR Policy Candidate Path Descriptor Format

Where:

* Type: 554

* Length: variable (valid values are 24, 36 or 48 octets)

* Protocol-Origin: 1-octet field which identifies the protocol or component which is responsible for the instantiation of this path as specified in section 2.3 of [RFC9256]. The protocol-origin codepoints to be used are listed in Section 8.4.

* Flags: 1-octet field with following bit positions defined. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

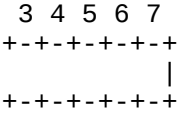


Where:

- E-Flag: Indicates the encoding of endpoint as IPv6 address when set and IPv4 address when clear

- O-Flag: Indicates the encoding of originator address as IPv6 address when set and IPv4 address when clear

* Reserved: 2 octets which MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Endpoint: 4 or 16 octets (as indicated by the flags) containing the address of the endpoint of the SR Policy as specified in section 2.1 of [RFC9256].

* Color: 4 octets that indicate the color of the SR Policy as specified in section 2.1 of [RFC9256].

* Originator ASN: 4 octets to carry the 4-byte encoding of the ASN of the originator. Refer to section 2.4 of [RFC9256] for details.

* Originator Address: 4 or 16 octets (as indicated by the flags) to carry the address of the originator. Refer to section 2.4 of [RFC9256] for details.

* Discriminator: 4 octets to carry the discriminator of the path. Refer to section 2.5 of [RFC9256] for details.

## 5. SR Policy State TLVs

This section defines the various TLVs which enable the headend to report the state at the SR Policy candidate path level. These TLVs (and their sub-TLVs) are carried in the optional non-transitive BGP-LS Attribute defined in [RFC9552] associated with the SR Policy Candidate Path NLRI type.

The detailed procedures for the advertisement are described in Section 6.

### 5.1. SR Binding SID TLV

The SR Binding SID (BSID) is an optional TLV that is used to report the BSID and its attributes for the SR Policy candidate path. The TLV MAY also optionally contain the Specified BSID value for reporting as described in section 6.2.3 of [RFC9256]. Only a single instance of this TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

The TLV has the following format:

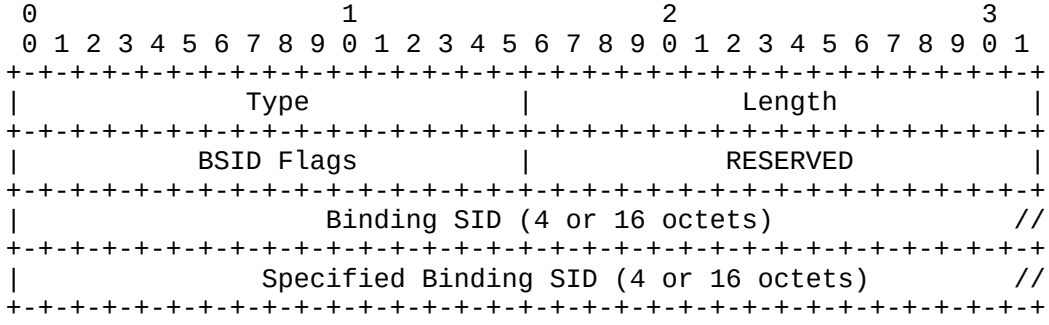


Figure 5 SR Binding SID TLV Format

Where:

* Type: 1201

* Length: variable (valid values are 12 or 36 octets)

* BSID Flags: 2-octet field that indicates attribute and status of the Binding SID (BSID) associated with this candidate path. The following bit positions are defined and the semantics are described in detail in section 6.2 of [RFC9256]. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

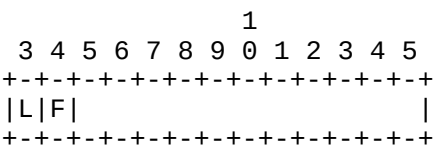


Where:

- D-Flag: Indicates the dataplane for the BSIDs and if they are 16 octet SRv6 SID (when set) or are 4 octet SR/MPLS label value (when clear).

- B-Flag: Indicates the allocation of the value in the BSID field when set and indicates that BSID is not allocated when clear.

- U-Flag: Indicates the specified BSID value is unavailable when set. When clear it indicates that this candidate path is using the specified BSID. This flag is ignored when there is no specified BSID.

- L-Flag: Indicates the BSID value is from the Segment Routing Local Block (SRLB) of the headend node when set and is from the local dynamic label pool when clear.

- F-Flag: Indicates the BSID value is one allocated from dynamic label pool due to fallback (e.g. when specified BSID is unavailable) when set and indicates that there has been no fallback for BSID allocation when clear.

* RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Binding SID: It indicates the operational or allocated BSID value based on the status flags.

* Specified BSID: It is used to report the explicitly specified BSID value regardless of whether it is successfully allocated or not. The field is set to value 0 when BSID has not been specified.

The BSID fields above depend on the dataplane (SRv6 or MPLS) indicated by the D-Flag. If D-Flag set (SRv6 dataplane), then the length of the BSID fields is 16 octets. If the D-Flag is clear (MPLS dataplane), then the length of the BSID fields is 4 octets. When carrying the MPLS Label, as shown in the figure below, the TC, S, and TTL (total of 12 bits) are RESERVED and MUST be set to 0 by the originator and MUST be ignored by a receiver.

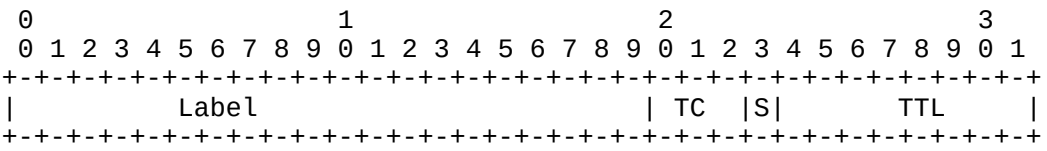


Figure 6 SR Binding SID Label Format

In the case of an SRv6, the SR Binding SID sub-TLV does not have the ability to signal the SRv6 Endpoint Behavior [RFC8986] or the structure of the SID. Therefore, the SR Binding SID sub-TLV SHOULD NOT be used for the advertisement of an SRv6 Binding SID. Instead, the SRv6 Binding SID TLV defined in Section 5.2 SHOULD be used for signaling of an SRv6 Binding SID. The use of the SR Binding SID sub-TLV for advertisement of the SRv6 Binding SID has been deprecated, and is documented here only for backward compatibility with implementations that followed early versions of this specification.

### 5.2. SRv6 Binding SID TLV

The SRv6 Binding SID (BSID) is an optional TLV that is used to report the SRv6 BSID and its attributes for the SR Policy candidate path. The TLV MAY also optionally contain the Specified SRv6 BSID value for reporting as described in section 6.2.3 of [RFC9256]. Multiple instances of this TLV may be used to report each of the SRv6 BSIDs associated with the candidate path.

The TLV has the following format:

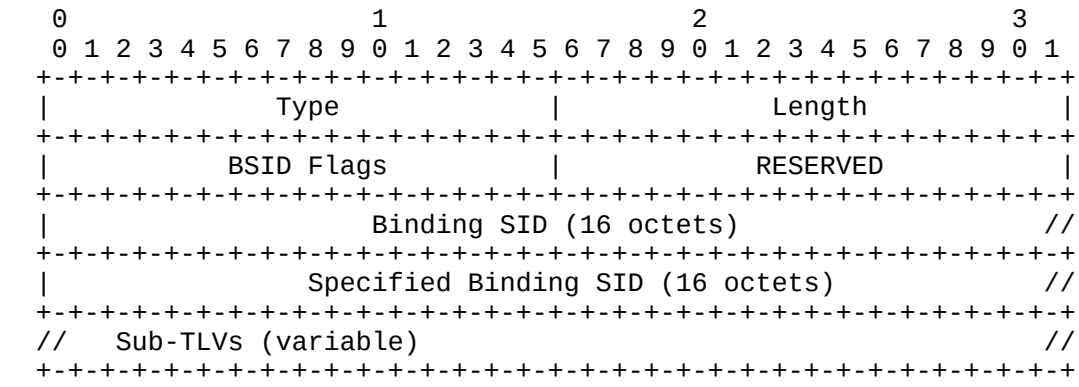


Figure 7 SRv6 Binding SID TLV Format

Where:

* Type: 1212

* Length: variable

* BSID Flags: 2-octet field that indicates attribute and status of the Binding SID (BSID) associated with this candidate path. The following bit positions are defined and the semantics are described in detail in section 6.2 of [RFC9256]. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

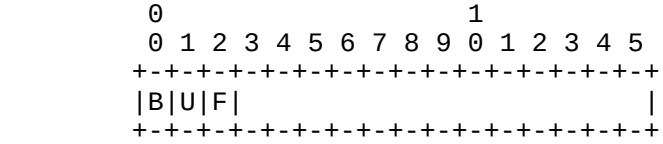


Where:

- B-Flag: Indicates the allocation of the value in the BSID field when set and indicates that BSID is not allocated when clear.

- U-Flag: Indicates the specified BSID value is unavailable when set. When clear it indicates that this candidate path is using the specified BSID. This flag is ignored when there is no specified BSID.

- F-Flag: Indicates the BSID value is one allocated dynamically due to fallback (e.g. when specified BSID is unavailable) when set and indicates that there has been no fallback for BSID allocation when clear.

* RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Binding SID: It indicates the operational or allocated BSID value based on the status flags.

* Specified BSID: It is used to report the explicitly specified BSID value regardless of whether it is successfully allocated or not. The field is set to value 0 when BSID has not been specified.

* Sub-TLVs: variable and contains any other optional attributes associated with the SRv6 BSID.

The SRv6 Endpoint Behavior TLV (1250) and the SRv6 SID Structure TLV (1252) MAY optionally be used as sub-TLVs of the SRv6 Binding SID TLV to indicate the SRv6 Endpoint behavior and SID structure for the Binding SID value in the TLV. [RFC9514] defines SRv6 Endpoint Behavior TLV And SRv6 SID Structure TLV.

### 5.3. SR Candidate Path State TLV

The SR Candidate Path State TLV provides the operational status and attributes of the SR Policy at the candidate path level. Only a single instance of this TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

The TLV has the following format:

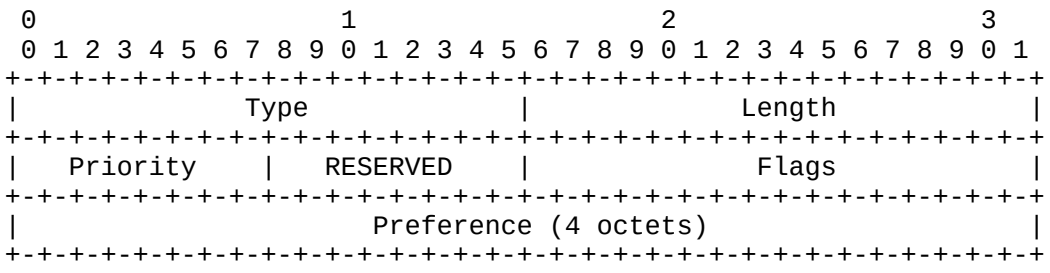


Figure 8 SR Candidate Path State TLV Format

Where:

* Type: 1202

* Length: 8 octets

* Priority: 1-octet value which indicates the priority of the candidate path. Refer Section 2.12 of [RFC9256].

* RESERVED: 1 octet. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Flags: 2-octet field that indicates attribute and status of the candidate path. The following bit positions are defined and the semantics are described in section 5 of [RFC9256] unless stated otherwise for individual flags. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

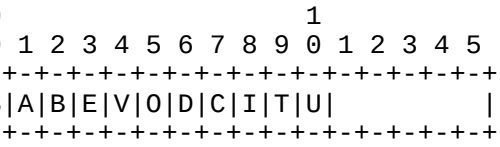


Where:

- S-Flag: Indicates the candidate path is in an administrative shut state when set and not in administrative shut state when clear.

- A-Flag: Indicates the candidate path is the active path (i.e. one provisioned in the forwarding plane as specified in section 2.9 of [RFC9256]) for the SR Policy when set and not the active path when clear.

- B-Flag: Indicates the candidate path is the backup path (i.e. one identified for path protection of the active path as specified in section 9.3 of [RFC9256]) for the SR Policy when set and not the backup path when clear.

- E-Flag: Indicates that the candidate path has been evaluated for validity (e.g. headend may evaluate candidate paths based on their preferences) when set and has not been evaluated for validity when clear.

- V-Flag: Indicates the candidate path has at least one valid SID-List when set and indicates no valid SID-List is available or evaluated when clear. When the E-Flag is clear (i.e. the candidate path has not been evaluated), then this flag MUST be set to 0 by the originator and ignored by the receiver.

- O-Flag: Indicates the candidate path was instantiated by the headend due to an on-demand nexthop trigger based on a local template when set and that the candidate path has not been instantiated due to on-demand nexthop trigger when clear. Refer to section 8.5 of [RFC9256] for details.

- D-Flag: Indicates the candidate path was delegated for computation to a PCE/controller when set and indicates that the candidate path has not been delegated for computation when clear.

- C-Flag: Indicates the candidate path was provisioned by a PCE/ controller when set and indicates that the candidate path was not provisioned by a PCE/controller when clear.

- I-Flag: Indicates the candidate path is to perform the "drop upon invalid" behavior when no other valid candidate path is available for this SR Policy when the flag is set. Refer to section 8.2 of [RFC9256] for details. When clear, it indicates that the candidate path is not enabled for the "drop upon invalid" behavior.

- T-Flag: Indicates the candidate path has been marked as eligible for use as a transit policy on the headend when set and not eligible for use as a transit policy when clear. Transit policy is a policy whose BSID can be used in the segment list of another SR Policy. Refer to section 8.3 of [RFC9256] for steering into a transit policy using its BSID.

- U-Flag: Indicates that this candidate path is reported as active and is dropping traffic as a result of the "drop upon invalid" behavior being activated for the SR Policy when set.

When clear, it indicates that the candidate path is not dropping traffic as a result of the "drop upon invalid" behavior. Refer to section 8.2 of [RFC9256] for details.

* Preference: 4-octet value which indicates the preference of the candidate path. Refer to section 2.7 of [RFC9256] for details.

### 5.4. SR Policy Name TLV

The SR Policy Name TLV is an optional TLV that is used to carry the symbolic name associated with the SR Policy. Only a single instance of this TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

The TLV has the following format:

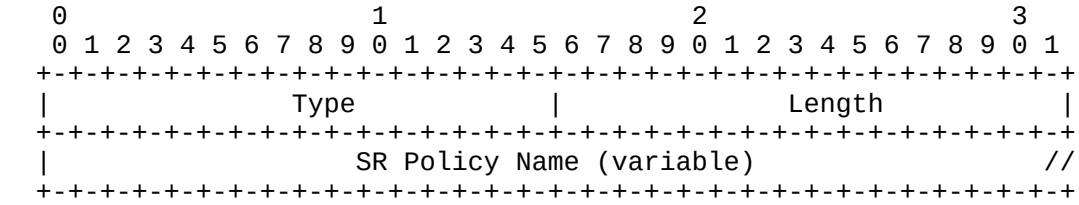


Figure 9 SR Policy Name TLV Format

Where:

* Type: 1213

* Length: variable

* SR Policy Name: Symbolic name for the SR Policy without a NULL terminator as specified in section 2.1 of [RFC9256]. It is RECOMMENDED that the size of the symbolic name be limited to 255 bytes. Implementations MAY choose to truncate long names to 255 bytes when signaling via BGP-LS.

### 5.5. SR Candidate Path Name TLV

The SR Candidate Path Name TLV is an optional TLV that is used to carry the symbolic name associated with the candidate path. Only a single instance of this TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

The TLV has the following format:

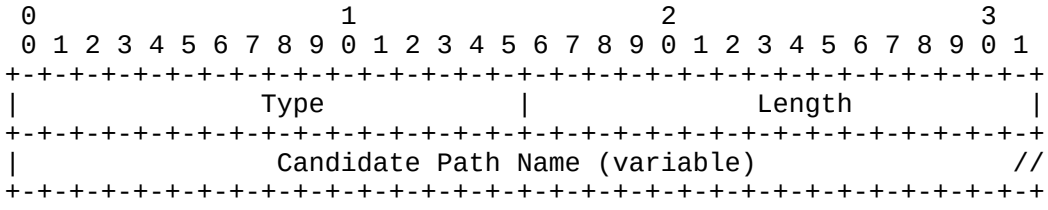


Figure 10  SR Candidate Path Name TLV Format

Where:

* Type: 1203

* Length: variable

* Candidate Path Name: Symbolic name for the SR Policy candidate path without a NULL terminator as specified in section 2.6 of [RFC9256]. It is RECOMMENDED that the size of the symbolic name be limited to 255 bytes. Implementations MAY choose to truncate long names to 255 bytes when signaling via BGP-LS.

## 5.6. SR Candidate Path Constraints TLV

The SR Candidate Path Constraints TLV is an optional TLV that is used to report the constraints associated with the candidate path. The constraints are generally applied to a dynamic candidate path which is computed either by the headend or may be delegated to a controller. The constraints may also be applied to an explicit path where the computation entity is expected to validate that the path satisfies the specified constraints and if not the path is to be invalidated (e.g., due to topology changes). Only a single instance of this TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

The TLV has the following format:

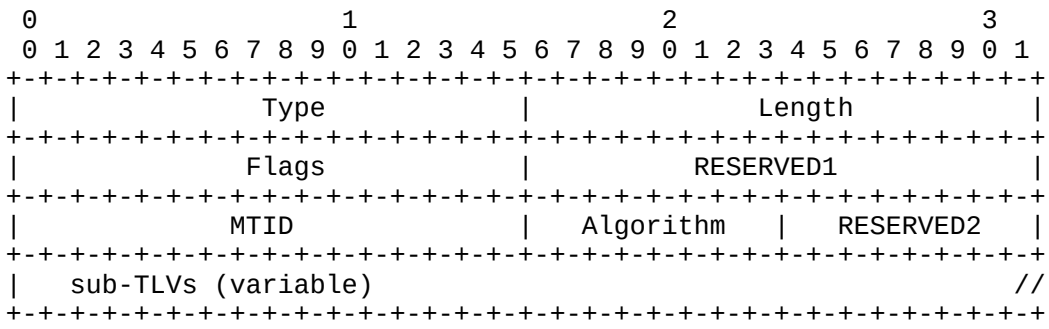


Figure 11 SR Candidate Path Constraints TLV Format

Where:

* Type: 1204

* Length: variable

* Flags: 2-octet field that indicates the constraints that are being applied to the candidate path. The following bit positions are defined and the other bits MUST be cleared by the originator and MUST be ignored by a receiver.

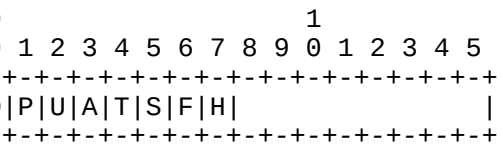


Where:

- D-Flag: Indicates that the candidate path uses SRv6 dataplane when set and SR/MPLS dataplane when clear

- P-Flag: Indicates that the candidate path prefers the use of only protected SIDs when set and indicates that the candidate path does not prefer the use of only protected SIDs when clear. This flag is mutually exclusive with the U-Flag (i.e., both these flags cannot be set at the same time).

- U-Flag: Indicates that the candidate path prefers the use of only unprotected SIDs when set and indicates that the candidate path does not prefer the use of only unprotected SIDs when clear. This flag is mutually exclusive with the P-Flag (i.e., both these flags cannot be set at the same time).

- A-Flag: Indicates that the candidate path uses only the SIDs belonging to the specified SR Algorithm when set and indicates that the candidate path does not use only the SIDs belonging to the specified SR Algorithm when clear.

- T-Flag: Indicates that the candidate path uses only the SIDs belonging to the specified topology when set and indicates that the candidate path does not use only the SIDs belonging to the specified topology when clear.

- S-Flag: Indicates that the use of protected (P-Flag) or unprotected (U-Flag) SIDs becomes a strict constraint instead of a preference when set and indicates that there is no strict constraint (and only a preference) when clear.

- F-Flag: Indicates that the candidate path is fixed once computed and not modified except on operator intervention and indicates that the candidate path may be modified as part of recomputation when clear.

- H-Flag: Indicates that the candidate path uses only adjacency SIDs and traverses hop-by-hop over the links corresponding to those adjacency SIDs when set and indicates that the candidate path is not restricted to using only hop-by-hop adjacency SIDs when clear.

* RESERVED1: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* MTID: Indicates the multi-topology identifier of the IGP topology that is preferred to be used when the path is set up. When the T-flag is set then the path is strictly using the specified topology SIDs only.

* Algorithm: Indicates the algorithm that is preferred to be used when the path is set up. When the A-flag is set then the path is strictly using the specified algorithm SIDs only. The algorithm values are from IGP Algorithm Types registry under the IANA Interior Gateway Protocol (IGP) Parameters.

* RESERVED2: 1 octet. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* sub-TLVs: one or more optional sub-TLVs MAY be included in this TLV to describe other constraints. These sub-TLVs are: SR Affinity Constraint, SR SRLG Constraint, SR Bandwidth Constraint, SR Disjoint Group Constraint, SR Bidirectional Group Constraint, and SR Metric Constraint.

These constraint sub-TLVs are defined below.

### 5.6.1. SR Affinity Constraint Sub-TLV

The SR Affinity Constraint sub-TLV is an optional sub-TLV of the SR Candidate Path Constraints TLV that is used to carry the affinity constraints [RFC2702] associated with the candidate path. The affinity is expressed in terms of Extended Admin Group (EAG) as defined in [RFC7308]. Only a single instance of this sub-TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

The sub-TLV has the following format:

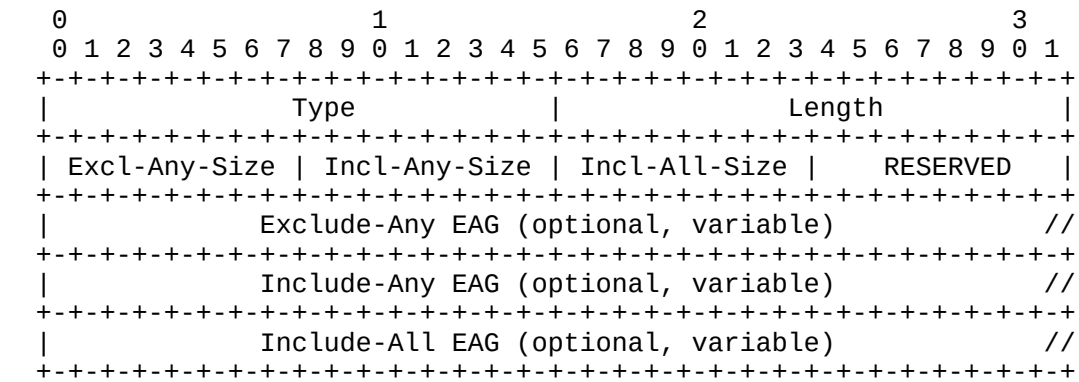


Figure 12 SR Affinity Constraints Sub-TLV Format

Where:

* Type: 1208

* Length: variable, dependent on the size of the Extended Admin Group. MUST be a non-zero multiple of 4 octets.

* Exclude-Any-Size: one octet to indicate the size of Exclude-Any EAG bitmask size in multiples of 4 octets. (e.g. value 0 indicates the Exclude-Any EAG field is skipped, value 1 indicates that 4 octets of Exclude-Any EAG is included)

* Include-Any-Size: one octet to indicate the size of Include-Any EAG bitmask size in multiples of 4 octets. (e.g. value 0 indicates the Include-Any EAG field is skipped, value 1 indicates that 4 octets of Include-Any EAG is included)

* Include-All-Size: one octet to indicate the size of Include-All EAG bitmask size in multiples of 4 octets. (e.g. value 0 indicates the Include-All EAG field is skipped, value 1 indicates that 4 octets of Include-All EAG is included)

* RESERVED: 1 octet. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Exclude-Any EAG: the bitmask used to represent the affinities that have been excluded from the path.

* Include-Any EAG: the bitmask used to represent the affinities that have been included in the path.

* Include-All EAG: the bitmask used to represent all the affinities that have been included in the path.

### 5.6.2. SR SRLG Constraint Sub-TLV

The SR SRLG Constraint sub-TLV is an optional sub-TLV of the SR Candidate Path Constraints TLV that is used to carry the Shared Risk Link Group (SRLG) values [RFC4202] that have been excluded from the candidate path. Only a single instance of this sub-TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

The sub-TLV has the following format:

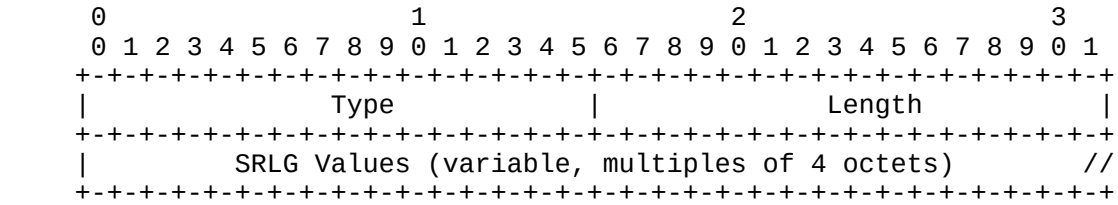


Figure 13 SR SRLG Constraints Sub-TLV Format

Where:

* Type: 1209

* Length: variable, dependent on the number of SRLGs encoded. MUST be a non-zero multiple of 4 octets.

* SRLG Values: One or more SRLG values. Each SRLG value is of 4 octets.

### 5.6.3. SR Bandwidth Constraint Sub-TLV

The SR Bandwidth Constraint sub-TLV is an optional sub-TLV of the SR Candidate Path Constraints TLV that is used to indicate the bandwidth that has been requested for the candidate path. Only a single instance of this sub-TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

The sub-TLV has the following format:

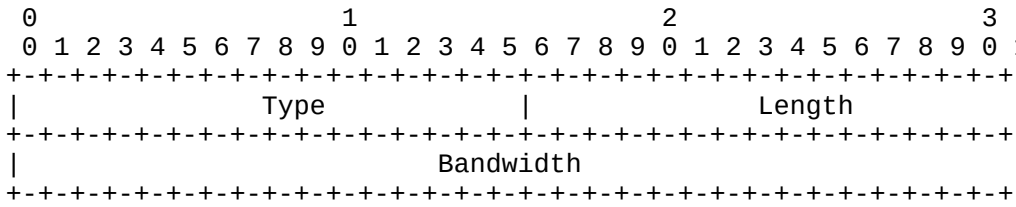


Figure 14 SR Bandwidth Constraints Sub-TLV Format

Where:

* Type: 1210
* Length: 4 octets
* Bandwidth: 4 octets which specify the desired bandwidth in unit of bytes per second in IEEE floating point format [IEEE754].

### 5.6.4. SR Disjoint Group Constraint Sub-TLV

The SR Disjoint Group Constraint sub-TLV is an optional sub-TLV of the SR Candidate Path Constraints TLV that is used to carry the disjointness constraint associated with the candidate path. The disjointness between two SR Policy Candidate Paths is expressed by associating them with the same disjoint group identifier and then specifying the type of disjointness required between their paths. The types of disjointness are described in section 3 of [RFC8800] where the level of disjointness increases in the order: link, node, SRLG, Node + SRLG. The computation is expected to achieve the highest level of disjointness requested and when that is not possible then fall back to a lesser level progressively based on the levels indicated. Only a single instance of this sub-TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

The sub-TLV has the following format:

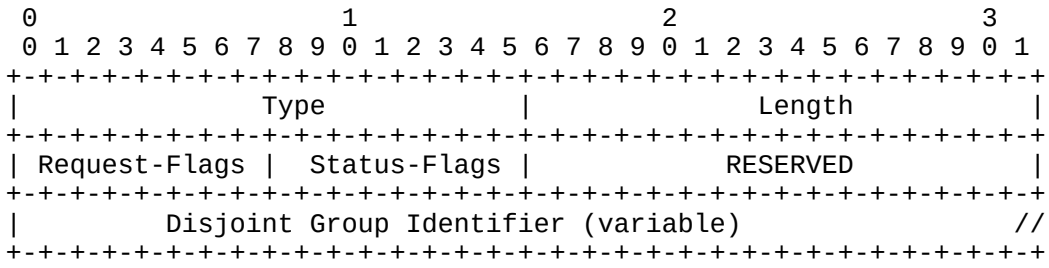


Figure 15 SR Disjoint Group Constraints Sub-TLV Format

Where:

* Type: 1211

* Length: Variable. Minimum of 8 octets.

* Request Flags: one octet to indicate the level of disjointness requested as specified in the form of flags. The following flags are defined and the other bits MUST be cleared by the originator and MUST be ignored by a receiver.

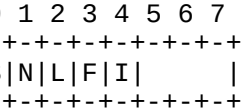


Where:

- S-Flag: Indicates that SRLG disjointness is requested when set and indicates that SRLG disjointness is not requested when clear.

- N-Flag: Indicates that node disjointness is requested when set and indicates that node disjointness is not requested when clear.

- L-Flag: Indicates that link disjointness is requested when set and indicates that the link disjointness is not requested when clear.

- F-Flag: Indicates that the computation may fall back to a lower level of disjointness amongst the ones requested when all cannot be achieved when set and indicates that fallback to a lower level of disjointness is not allowed when clear.

- I-Flag: Indicates that the computation may fall back to the default best path (e.g. IGP path) in case of none of the desired disjointness can be achieved when set and indicates that fallback to the default best path is not allowed when clear.

* Status Flags: one octet to indicate the level of disjointness that has been achieved by the computation as specified in the form of flags. The following flags are defined and the other bits MUST be cleared by the originator and MUST be ignored by a receiver.


Where:

- S-Flag: Indicates that SRLG disjointness is achieved when set and indicates that SRLG disjointness is not achieved when clear.

- N-Flag: Indicates that node disjointness is achieved when set and indicates that node disjointness was not achieved when clear.

- L-Flag: Indicates that link disjointness is achieved when set and indicates that link disjointness was not achieved when clear.

- F-Flag: Indicates that the computation has fallen back to a lower level of disjointness than requested when set and indicates that there has been no fallback to a lower level of disjointness when clear.

- I-Flag: Indicates that the computation has fallen back to the best path (e.g. IGP path) and disjointness has not been achieved when set and indicates that there has been no fallback to best path when clear.

- X-Flag : Indicates that the disjointness constraint could not be achieved and hence path has been invalidated when set and indicates that the path has not been invalidated due to unmet disjointness constraints when clear.

* RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Disjoint Group Identifier: 4-octet value that is the group identifier for a set of disjoint paths. Alternatively, this field MAY contain the entire PCEP Association Object as specified in section 6.1 of [RFC8697] (including its optional TLVs) when PCEP is used for the signaling the SR Policy candidate path and where the BGP-LS Producer is unable to determine the group identifier that can be accommodated in a 4-octet value (since PCEP supports multiple methods of encoding an association identifier). Note that the parsing of the PCEP object is expected to be performed only by the BGP-LS Consumer (hence, outside the scope of this document) and not by any BGP Speaker as specified in [RFC9552]. If the PCEP object size is such that the update for a single SR Policy Candidate Path NLRI would exceed the supported BGP message size by the implementation, then the PCEP Association Object MUST NOT be encoded and this sub-TLV skipped along with an error log. Refer section 5.3 of [RFC9552] for discussion on implications of encoding large sets of information into BGP-LS.

#### 5.6.5. SR Bidirectional Group Constraint Sub-TLV

The SR Bidirectional Group Constraint sub-TLV is an optional sub-TLV of the SR Candidate Path Constraints TLV that is used to carry the bidirectional constraint associated with the candidate path. The bidirectional relationship between two SR Policy Candidate Paths is expressed by associating them with the same bidirectional group identifier and then specifying the type of bidirectional routing required between their paths. Only a single instance of this sub-TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

The sub-TLV has the following format:

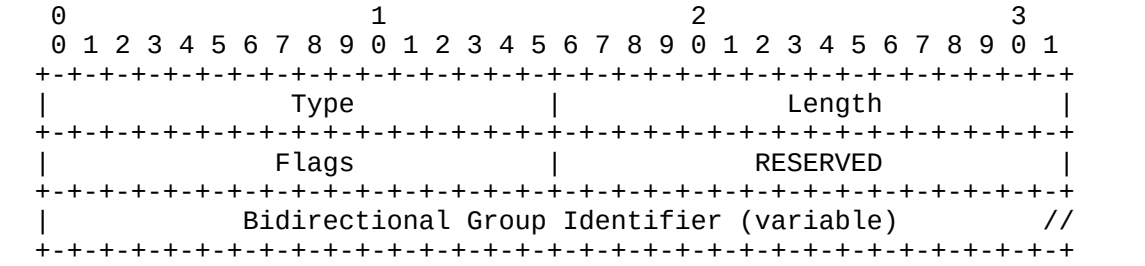


Figure 16 SR Bidirectional Group Constraints Sub-TLV Format

Where:

* Type: 1214

* Length: Variable. Minimum of 8 octets.

* Flags: two octets to indicate the bidirectional path setup information as specified in the form of flags. The following flags are defined and the other bits MUST be cleared by the originator and MUST be ignored by a receiver.

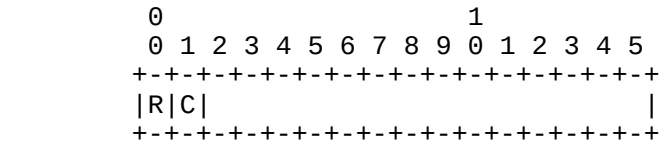


  Where:

  - R-Flag: Indicates that this candidate path of the SR Policy forms the reverse path when the R-Flag is set. If the R-Flag is clear, this candidate path forms the forward path.

  - C-Flag: Indicates that the bidirectional path is co-routed when set and indicates that the bidirectional path is not co-routed when clear.

* RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Bidirectional Group Identifier: 4-octet value that is the group identifier for a set of bidirectional paths. Alternatively, this field MAY contain the entire PCEP Association Object as specified in section 6.1 of [RFC8697] (including its optional TLVs) when PCEP is used for the signaling the SR Policy candidate path and where the BGP-LS Producer is unable to determine the group identifier that can be accommodated in a 4-octet value (since PCEP supports multiple methods of encoding an association identifier). Note that the parsing of the PCEP object is expected to be performed only by the BGP-LS Consumer (hence, outside the scope of this document) and not by any BGP Speaker as specified in [RFC9552]. If the PCEP object size is such that the update for a single SR Policy Candidate Path NLRI would exceed the supported BGP message size by the implementation, then the PCEP Association Object MUST NOT be encoded and this sub-TLV skipped along with an error log. Refer section 5.3 of [RFC9552] for discussion on implications of encoding large sets of information into BGP-LS.

### 5.6.6. SR Metric Constraint Sub-TLV

The SR Metric Constraint sub-TLV is an optional sub-TLV of the SR Candidate Path Constraints TLV that is used to report the optimization metric of the candidate path. For a dynamic path computation, it is used to report the optimization metric used along with its parameters. For an explicit path, this sub-TLV MAY be used to report the metric margin or bound to be used for validation (i.e., the path is invalidated if the metric is beyond specified values). Multiple instances of this sub-TLV may be used to report different metric type uses.

The sub-TLV has the following format:

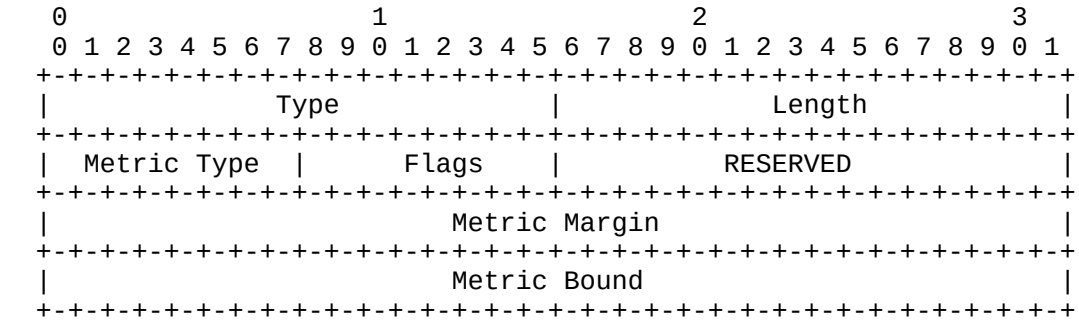


Figure 17 SR Metric Constraints Sub-TLV Format

Where:

* Type: 1215

* Length: 12 octets

* Metric Type: 1-octet field which identifies the type of the metric being used. The Table 1 below lists the metric types introduced by this document along with reference for each. Where the references are for IS-IS and OSPF specifications, those metric types are defined for a link while in the SR Policy context those relate to the candidate path or the segment list. The metric type code points that may be used in this sub-TLV are also listed in Section 8.6 of this document. Note that the metric type in this field is not taken from the "IGP Metric Type" registry from IANA "IGP Parameters" and is a separate registry that includes IGP Metric Types as well as metric types specific to SR Policy path computation. Additional metric types may be introduced by future documents. This document does not make any assumption of a smaller metric value being better than a higher metric value; that

is something dependent on the semantics of the specific metric type. The document uses the words "best" and "worst" to abstract this aspect when referring to metric margins and bounds.

- Type 0: IGP: In IS-IS, this is known as the default metric and specified in section 3 of [RFC5305]. This is known as metric in both OSPFv2 [RFC2328] and OSPFv3 [RFC5340].

- Type 1: Min Unidirectional Delay: This is specified in section 4.2 of [RFC8570] for IS-IS and in section 4.2 of [RFC7471] for OSPFv2/OSPFv3.

- Type 2: TE: This is specified in section 3.7 of [RFC5305] as the TE default metric for IS-IS, in section 2.5.5 of [RFC3630] for OSPFv2, and in section 4 of [RFC5329] for OSPFv3.

- Type 3: Hop Count: This is specified in section 7.8 of [RFC5440].

- Type 4: SID List Length: This is specified in section 4.5 of [RFC8664].

- Type 5: Bandwidth: This is specified in section 4 of [I-D.ietf-lsr-flex-algo-bw-con].

- Type 6: Average Unidirectional Delay: This is specified in section 4.1 of [RFC8570] for IS-IS and in section 4.1 of [RFC7471] for OSPFv2/OSPFv3.

- Type 7: Unidirectional Delay Variation: This is specified in section 4.3 of [RFC8570] for IS-IS and in section 4.3 of [RFC7471] for OSPFv2/OSPFv3.

- Type 8: Loss: This is specified in section 4.4 of [RFC8570] for IS-IS and in section 4.4 of [RFC7471] for OSPFv2/OSPFv3.

- Types 128 to 255 (both inclusive): User Defined: This is specified for IS-IS and OSPF in section 2 of [I-D.ietf-lsr-flex-algo-bw-con].

* Flags: 1-octet field that indicates the validity of the metric fields and their semantics. The following bit positions are defined and the other bits MUST be cleared by the originator and MUST be ignored by a receiver.

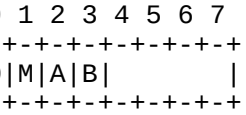


Where:

- O-Flag: Indicates that this is the optimization metric being reported for a dynamic candidate path when set and indicates that the metric is not the optimization metric when clear. This bit MUST NOT be set in more than one instance of this TLV for a given candidate path advertisement.

- M-Flag: Indicates that the metric margin allowed is specified when set and indicates that metric margin allowed is not specified when clear.

- A-Flag: Indicates that the metric margin is specified as an absolute value when set and is expressed as a percentage of the minimum metric when clear.

- B-Flag: Indicates that the metric bound allowed for the path is specified when set and indicates that metric bound is not specified when clear.

* RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Metric Margin: 4-octet value which indicates the metric margin when the M-flag is set. The metric margin is specified, depending on the A-flag, as either an absolute value or as a percentage of the best computed path metric based on the specified constraints for path calculation. The metric margin allows for the metric value of the computed path to vary (depending on the semantics of the specific metric type) from the best metric value possible to optimize for other factors (that are not specified as constraints) such as bandwidth availability, minimal SID stack depth, and maximizing of ECMP for the SR path computed.

* Metric Bound: 4-octet value which indicates the worst metric value (depending on the semantics of the specific metric type) that is allowed when the B-flag is set. If the computed path metric crosses the specified bound value then the path is considered invalid.

The absolute metric margin and the metric bound values are encoded as specified for each metric type. For metric types that are smaller than 4 octets in size, the most significant bits are filled with zeros. The percentage metric margin is encoded as an unsigned integer percentage value.

## 5.7. SR Segment List TLV

The SR Segment List TLV is used to report a single SID-List of a candidate path. Multiple instances of this TLV may be used to report multiple SID-Lists of a candidate path.

The TLV has the following format:

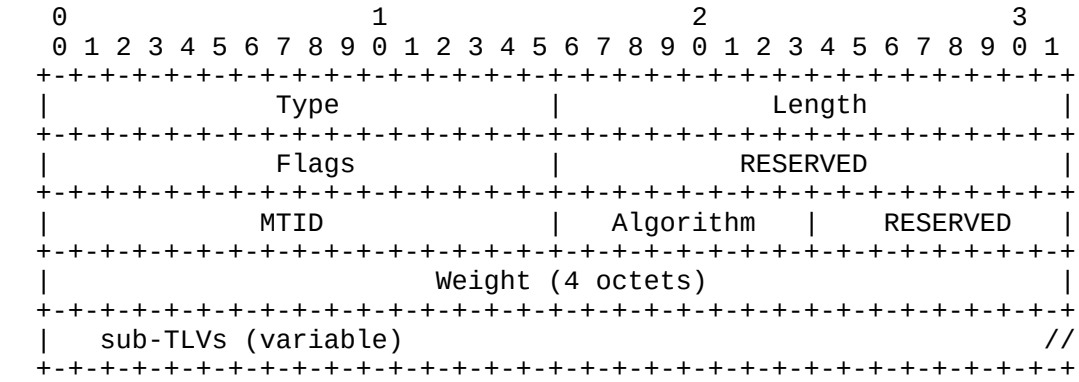


Figure 18 SR Segment List TLV Format

Where:

* Type: 1205

* Length: variable

* Flags: 2-octet field that indicates attribute and status of the SID-List.The following bit positions are defined and the semantics are described in detail in [RFC9256]. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

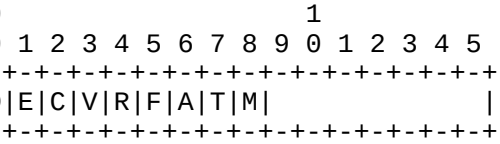


Where:

- D-Flag: Indicates the SID-List consists of SRv6 SIDs when set and indicates it consists of SR/MPLS labels when clear.

- E-Flag: Indicates that SID-List is associated with an explicit candidate path when set and with a dynamic candidate path when clear. All segment lists of a given candidate path MUST be either explicit or dynamic and in case of inconsistency, the receiver MAY consider them all to be dynamic.

- C-Flag: Indicates that SID-List has been computed for a dynamic path when set. It is always reported as set for explicit paths. When clear, it indicates that the SID-List has not been computed for a dynamic path.

- V-Flag: Indicates the SID-List has passed verification or its verification was not required when set and failed verification when clear.

- R-Flag: Indicates that the first Segment has been resolved when set and failed resolution when clear.

- F-Flag: Indicates that the computation for the dynamic path failed when set and succeeded (or not required in case of explicit path) when clear.

- A-Flag: Indicates that all the SIDs in the SID-List belong to the specified algorithm when set and indicates that not all the SIDs belong to the specified algorithm when clear.

- T-Flag: Indicates that all the SIDs in the SID-List belong to the specified topology (identified by the multi-topology ID) when set and indicates that not all the SIDs belong to the specified topology when clear.

- M-Flag: Indicates that the SID-list has been removed from the forwarding plane due to fault detection by a monitoring mechanism (e.g. BFD) when set and indicates no fault detected or monitoring is not being done when clear.

* RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* MTID: 2 octets that indicates the multi-topology identifier of the IGP topology that is to be used when the T-flag is set.

* Algorithm: 1 octet that indicates the algorithm of the SIDs used in the SID-List when the A-flag is set. The algorithm values are from IGP Algorithm Types registry under the IANA Interior Gateway Protocol (IGP) Parameters.

* RESERVED: 1 octet. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Weight: 4-octet field that indicates the weight associated with the SID-List for weighted load-balancing. Refer to section 2.2 and 2.11 of [RFC9256].

* Sub-TLVs: variable and contains the ordered set of Segments and any other optional attributes associated with the specific SID-List.

The SR Segment sub-TLV (defined in Section 5.7.1) MUST be included as an ordered set of sub-TLVs within the SR Segment List TLV when the SID-List is not empty. A SID-List may be empty in certain situations (e.g. for a dynamic path) where the headend has not yet performed the computation and hence not derived the segments required for the path. In such cases where the SID-LIST is empty, the SR Segment List TLV MUST NOT include any SR Segment sub-TLVs.

### 5.7.1. SR Segment Sub-TLV

The SR Segment sub-TLV describes a single segment in a SID-List. One or more instances of this sub-TLV in an ordered manner constitute a SID-List for an SR Policy candidate path. It is a sub-TLV of the SR Segment List TLV and it has the following format:

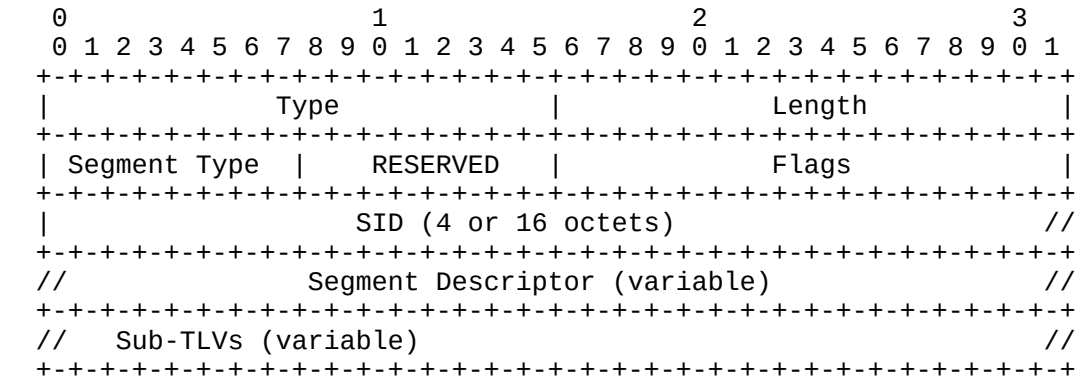


Figure 19 SR Segment Sub-TLV Format

Where:

* Type: 1206

* Length: variable

* Segment Type: 1 octet which indicates the type of segment. Initial values are specified by this document (see Section 5.7.1.1 for details). Additional segment types are possible, but out of scope for this document.

* RESERVED: 1 octet. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Flags: 2-octet field that indicates attribute and status of the Segment and its SID. The following bit positions are defined and the semantics are described in section 5 of [RFC9256]. Other bits MUST be cleared by the originator and MUST be ignored by a receiver.

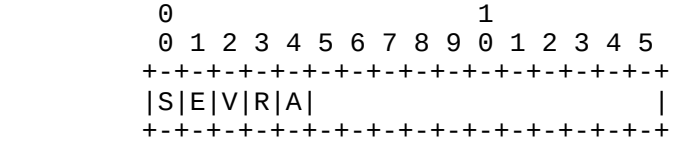


Where:

- S-Flag: Indicates the presence of SID value in the SID field when set and that no value is indicated when clear.

- E-Flag: Indicates the SID value is explicitly provisioned value (locally on headend or via controller/PCE) when set and is a dynamically resolved value by headend when clear

- V-Flag: Indicates the SID has passed verification or did not require verification when set. When V-Flag is clear, it indicates the SID has failed verification.

- R-Flag: Indicates the SID has been resolved or did not require resolution (e.g. because it is not the first SID) when set. When R-Flag is clear, it indicates the SID has failed resolution.

- A-Flag: Indicates that the Algorithm indicated in the Segment descriptor is valid when set. When clear, it indicates that the headend is unable to determine the algorithm of the SID.

* SID: 4 octets carrying the MPLS Label or 16 octets carrying the SRv6 SID based on the Segment Type. When carrying the MPLS Label, as shown in the figure below, the TC, S, and TTL (total of 12 bits) are RESERVED and MUST be set to 0 by the originator and MUST be ignored by a receiver.

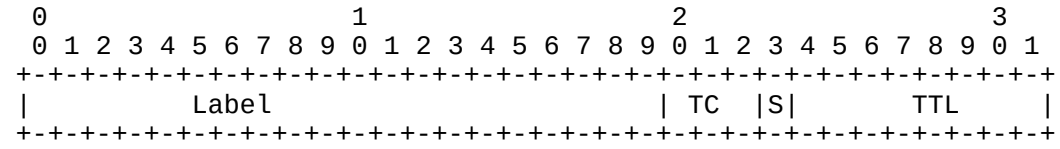


* Segment Descriptor: variable size Segment descriptor based on the type of segment (refer to Section 5.7.1.1 for details)

* Sub-Sub-TLVs: variable and contains any other optional attributes associated with the specific segment.

The SRv6 Endpoint Behavior TLV (1250) and the SRv6 SID Structure TLV (1252) defined in [RFC9514] are used as sub-sub-TLVs of the SR Segment sub-TLV. These two sub-sub-TLVs are used to optionally indicate the SRv6 Endpoint behavior and SID structure when advertising the SRv6 specific segment types.

#### 5.7.1.1. Segment Descriptors

Section 4 of [RFC9256] defines multiple types of segments and their description. This section defines the encoding of the Segment Descriptors for each of those Segment types to be used in the Segment sub-TLV described previously in Section 5.7.1.

The following types are currently defined and their mapping to the respective segment types defined in [RFC9256]:

| Type | Segment Description |
|---|---|
| 1 | (Type A) SR-MPLS Label |
| 2 | (Type B) SRv6 SID as IPv6 address |
| 3 | (Type C) SR-MPLS Prefix SID as IPv4 Node Address |
| 4 | (Type D) SR-MPLS Prefix SID as IPv6 Node Global Address |
| 5 | (Type E) SR-MPLS Adjacency SID as IPv4 Node Address & Local Interface ID |
| 6 | (Type F) SR-MPLS Adjacency SID as IPv4 Local & Remote Interface Addresses |
| 7 | (Type G) SR-MPLS Adjacency SID as pair of IPv6 Global Address & Interface ID for Local & Remote nodes |
| 8 | (Type H) SR-MPLS Adjacency SID as pair of IPv6 Global Addresses for the Local & Remote Interface |
| 9 | (Type I) SRv6 END SID as IPv6 Node Global Address |
| 10 | (Type J) SRv6 END.X SID as pair of IPv6 Global Address & Interface ID for Local & Remote nodes |
| 11 | (Type K) SRv6 END.X SID as pair of IPv6 Global Addresses for the Local & Remote Interface |

Table 1  SR Segment Types

#### 5.7.1.1.1.  Type 1: SR-MPLS Label (Type A)

The Segment is SR-MPLS type and is specified simply as the label. The format of its Segment Descriptor is as follows:

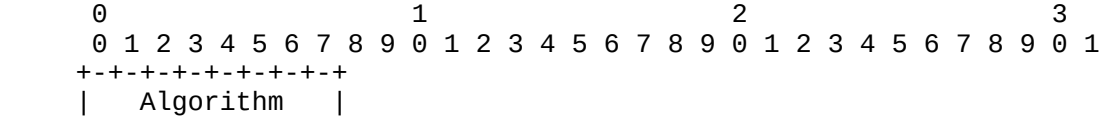


Figure 20  Type 1 Segment Descriptor

Where:

* Algorithm: 1-octet value that indicates the algorithm used for picking the SID.  This is valid only when the A-flag has been set in the Segment TLV.  The algorithm values are from IGP Algorithm Types registry under the IANA Interior Gateway Protocol (IGP) Parameters.

#### 5.7.1.1.2. Type 2: SRv6 SID (Type B)

The Segment is SRv6 type and is specified simply as the SRv6 SID address. The format of its Segment Descriptor is as follows:

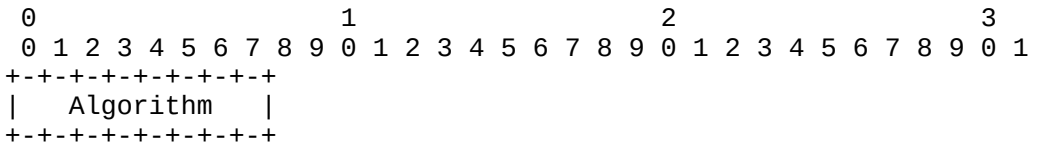


Figure 21 Type 2 Segment Descriptor

Where:

* Algorithm: 1-octet value that indicates the algorithm used for picking the SID. This is valid only when the A-flag has been set in the Segment TLV. The algorithm values are from IGP Algorithm Types registry under the IANA Interior Gateway Protocol (IGP) Parameters.

#### 5.7.1.1.3. Type 3: SR-MPLS Prefix SID for IPv4 (Type C)

The Segment is SR-MPLS Prefix SID type and is specified as an IPv4 node address. The format of its Segment Descriptor is as follows:

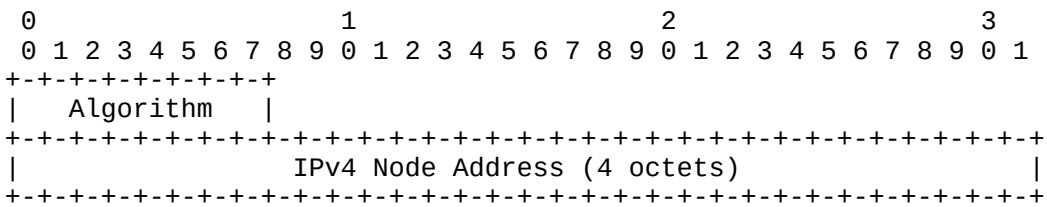


Figure 22 Type 3 Segment Descriptor

Where:

* Algorithm: 1-octet value that indicates the algorithm used for picking the SID. The algorithm values are from IGP Algorithm Types registry under the IANA Interior Gateway Protocol (IGP) Parameters.

* IPv4 Node Address: 4-octet value which carries the IPv4 address associated with the node

#### 5.7.1.1.4. Type 4: SR-MPLS Prefix SID for IPv6 (Type D)

The Segment is SR-MPLS Prefix SID type and is specified as an IPv6 global address. The format of its Segment Descriptor is as follows:

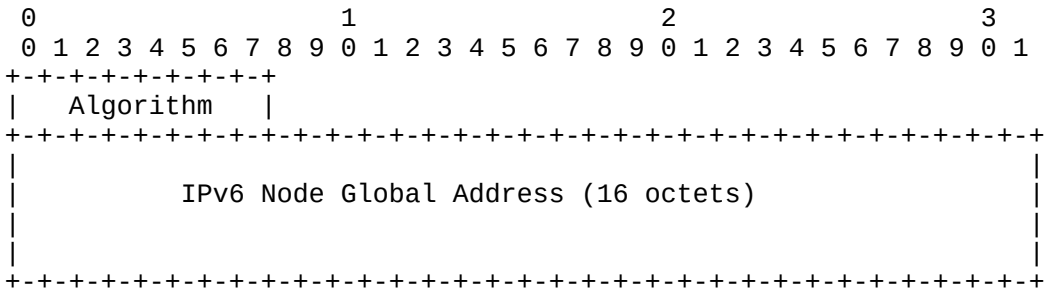


Figure 23 Type 4 Segment Descriptor

Where:

* Algorithm: 1-octet value that indicates the algorithm used for picking the SID. The algorithm values are from IGP Algorithm Types registry under the IANA Interior Gateway Protocol (IGP) Parameters.

* IPv6 Node Global Address: 16-octet value which carries the IPv6 global address associated with the node

#### 5.7.1.1.5. Type 5: SR-MPLS Adjacency SID for IPv4 with an Interface ID (Type E)

The Segment is SR-MPLS Adjacency SID type and is specified as an IPv4 node address along with the local interface ID on that node. The format of its Segment Descriptor is as follows:

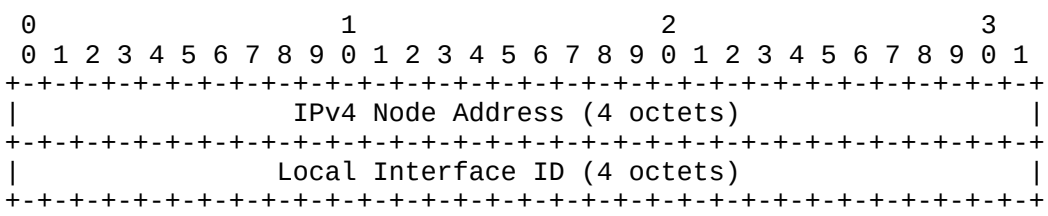


Figure 24 Type 5 Segment Descriptor

Where:

* IPv4 Node Address: 4-octet value which carries the IPv4 address associated with the node

* Local Interface ID: 4-octet value which carries the local interface ID of the node identified by the Node Address

#### 5.7.1.1.6. Type 6: SR-MPLS Adjacency SID for IPv4 with an Interface Address (Type F)

The Segment is SR-MPLS Adjacency SID type and is specified as a pair of IPv4 local and remote addresses. The format of its Segment Descriptor is as follows:

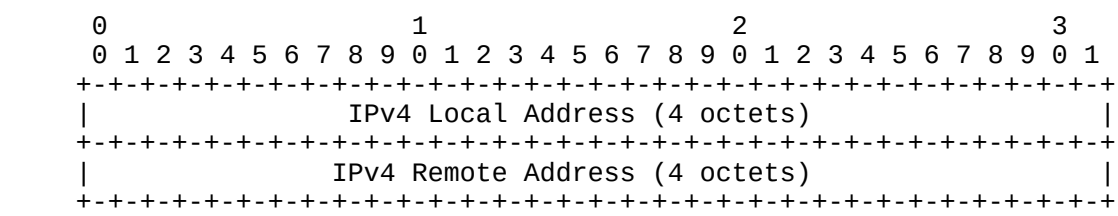


Figure 25 Type 6 Segment Descriptor

Where:

* IPv4 Local Address: 4-octet value which carries the local IPv4 address associated with the node’s interface

* IPv4 Remote Address: 4-octet value which carries the remote IPv4 address associated with interface on the node’s neighbor. This is optional and MAY be set to 0 when not used (e.g. when identifying point-to-point links).

#### 5.7.1.1.7. Type 7: SR-MPLS Adjacency SID for IPv6 with an interface ID (Type G)

The Segment is SR-MPLS Adjacency SID type and is specified as a pair of IPv6 global address and interface ID for local and remote nodes. The format of its Segment Descriptor is as follows:

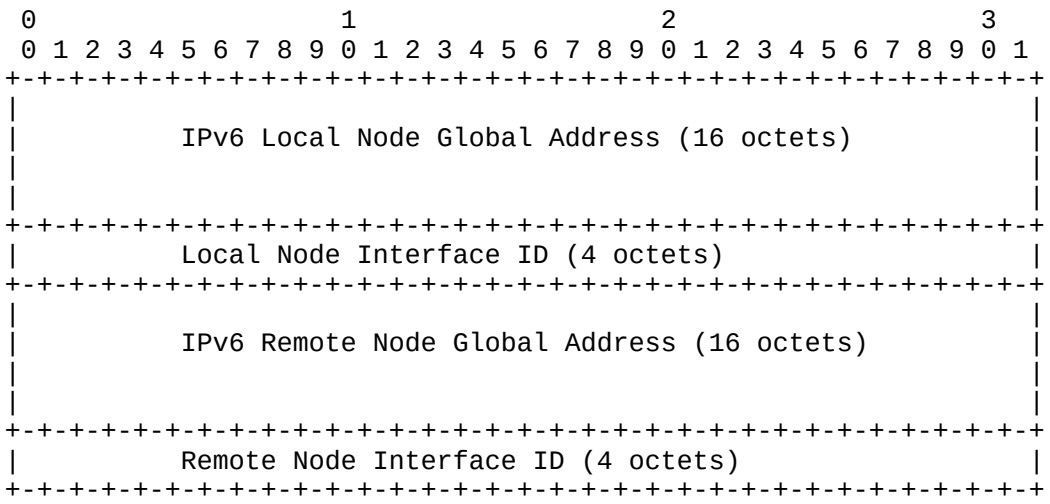


Figure 26 Type 7 Segment Descriptor

Where:

* IPv6 Local Node Global Address: 16-octet value which carries the IPv6 global address associated with the local node

* Local Node Interface ID : 4-octet value which carries the interface ID of the local node identified by the Local Node Address

* IPv6 Remote Node Global Address: 16-octet value which carries the IPv6 global address associated with the remote node. This is optional and MAY be set to 0 when not used (e.g. when identifying point-to-point links).

* Remote Node Interface ID: 4-octet value which carries the interface ID of the remote node identified by the Remote Node Address. This is optional and MAY be set to 0 when not used (e.g. when identifying point-to-point links).

#### 5.7.1.1.8. Type 8: SR-MPLS Adjacency SID for IPv6 with an Interface Address (Type H)

The Segment is SR-MPLS Adjacency SID type and is specified as a pair of IPv6 Global addresses for local and remote interface addresses. The format of its Segment Descriptor is as follows:

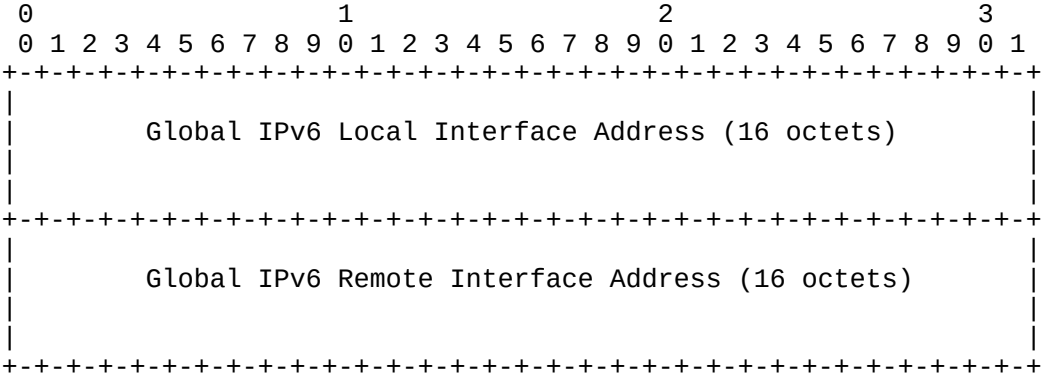


Figure 27 Type 8 Segment Descriptor

Where:

* IPv6 Local Address: 16-octet value which carries the local IPv6 address associated with the node’s interface

* IPv6 Remote Address: 16-octet value which carries the remote IPv6 address associated with the interface on the node’s neighbor

#### 5.7.1.1.9. Type 9: SRv6 END SID as IPv6 Node Address (Type I)

The Segment is SRv6 END SID type and is specified as an IPv6 global address. The format of its Segment Descriptor is as follows:

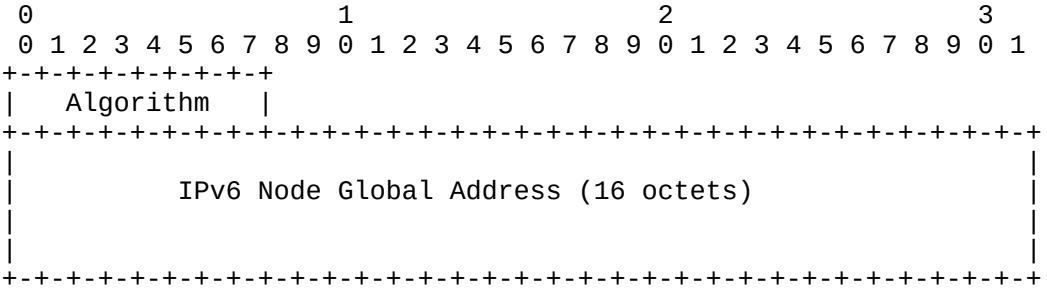


Figure 28 Type 9 Segment Descriptor

Where:

* Algorithm: 1-octet value that indicates the algorithm used for picking the SID. The algorithm values are from IGP Algorithm Types registry under the IANA Interior Gateway Protocol (IGP) Parameters.

* IPv6 Node Global Address: 16-octet value which carries the IPv6 global address associated with the node

#### 5.7.1.1.10. Type 10: SRv6 END.X SID as an Interface ID (Type J)

The Segment is SRv6 END.X SID type and is specified as a pair of IPv6 global address and interface ID for local and remote nodes. The format of its Segment Descriptor is as follows:

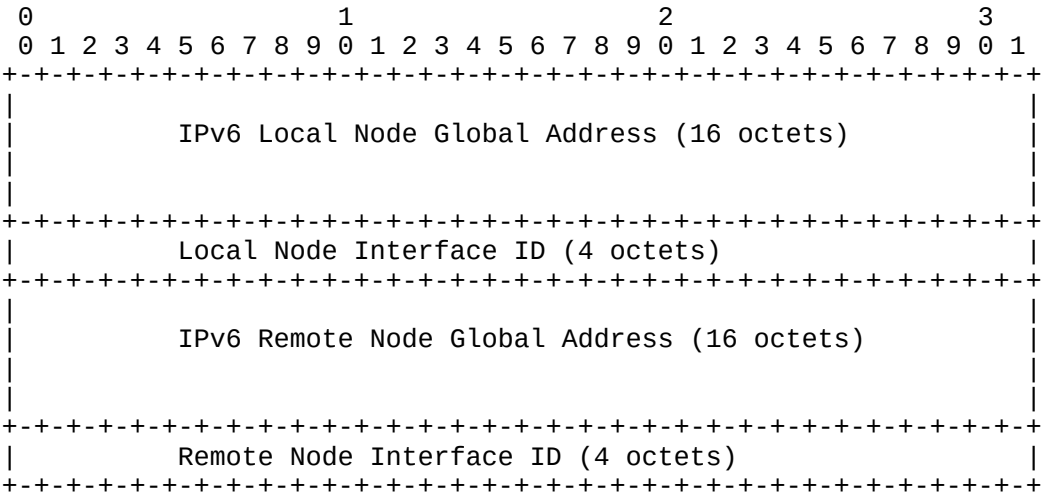


Figure 29 Type 10 Segment Descriptor

Where:

* IPv6 Local Node Global Address: 16-octet value which carries the IPv6 global address associated with the local node

* Local Node Interface ID: 4-octet value which carries the interface ID of the local node identified by the Local Node Address

* IPv6 Remote Node Global Address: 16-octet value which carries the IPv6 global address associated with the remote node. This is optional and MAY be set to 0 when not used (e.g. when identifying point-to-point links).

* Remote Node Interface ID: 4-octet value which carries the interface ID of the remote node identified by the Remote Node Address. This is optional and MAY be set to 0 when not used (e.g. when identifying point-to-point links).

#### 5.7.1.1.11. Type 11: SRv6 END.X SID as an Interface Address (Type K)

The Segment is SRv6 END.X SID type and is specified as a pair of IPv6 Global addresses for local and remote interface addresses. The format of its Segment Descriptor is as follows:

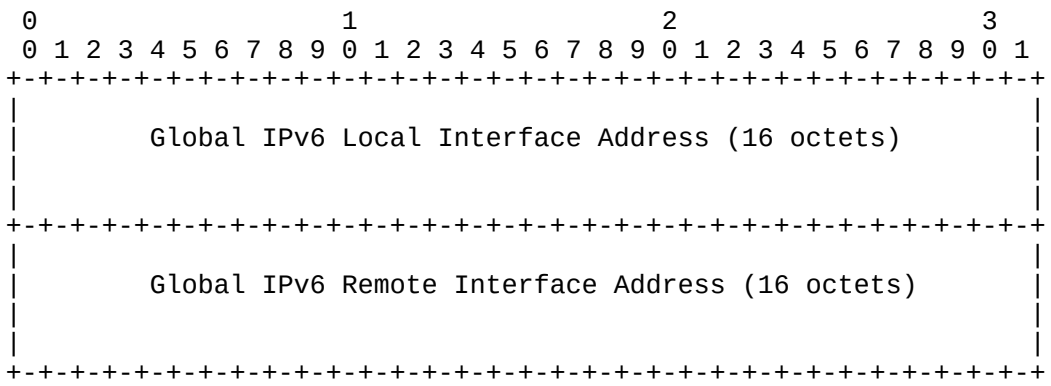


Figure 30 Type 11 Segment Descriptor

Where:

* IPv6 Local Address: 16-octet value which carries the local IPv6 address associated with the node’s interface

* IPv6 Remote Address: 16-octet value which carries the remote IPv6 address associated with the interface on the node’s neighbor

### 5.7.2. SR Segment List Metric Sub-TLV

The SR Segment List Metric sub-TLV reports the computed metric of the specific SID-List. It is used to report the type of metric and its computed value by the computation entity (i.e., either the headend or the controller when the path is delegated) when available. More than one instance of this sub-TLV may be present in SR Segment List to report metric values of different metric types. The metric margin and bound may be optionally reported using this sub-TLV when this information is not being reported using the SR Metric Constraint sub-TLV (refer to Section 5.6.6) at the SR candidate path level.

It is a sub-TLV of the SR Segment List TLV and has the following format:

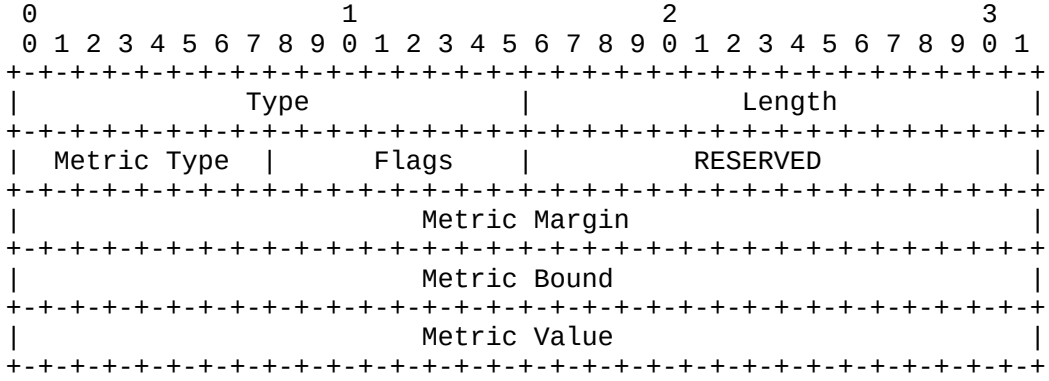


Figure 31  SR Segment List Metric Sub-TLV Format

Where:

* Type: 1207

* Length: 16 octets

* Metric Type: 1-octet field which identifies the type of metric. The semantics are the same as the Metric Type field in the SR Metric Constraints sub-TLV in Section 5.6.6 of this document.

* Flags: 1-octet field that indicates the validity of the metric fields and their semantics. The following bit positions are defined and the other bits MUST be cleared by the originator and MUST be ignored by a receiver.

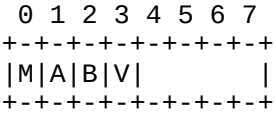


Where:

- M-Flag: Indicates that the metric margin allowed for this path computation is specified when set and indicates that metric margin allowed is not specified when clear.

- A-Flag: Indicates that the metric margin is specified as an absolute value when set and is expressed as a percentage of the minimum metric when clear.

  - B-Flag: Indicates that the metric bound allowed for the path is specified when set and indicates that metric bound is not specified when clear.

  - V-Flag: Indicates that the metric value computed is being reported when set and indicates that the computed metric value is not being reported when clear.

* RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.

* Metric Margin: 4-octet value which indicates the metric margin value when the M-flag is set. The metric margin is specified, depending on the A-flag, as either an absolute value or as a percentage of the best computed path metric based on the specified constraints for path calculation. The metric margin allows for the metric value of the computed path to vary (depending on the semantics of the specific metric type) from the best metric value possible to optimize for other factors (that are not specified as constraints) such as bandwidth availability, minimal SID stack depth, and maximizing of ECMP for the SR path computed.

* Metric Bound: 4-octet value which indicates the worst metric value (depending on the semantics of the specific metric type) that is allowed when the B-flag is set. If the computed path metric crosses the specified bound value then the path is considered invalid.

* Metric Value: 4-octet value which indicates the metric of the computed path when the V-flag is set. This value is available and reported when the computation is successful and a valid path is available.

The absolute metric margin, metric bound, and metric values are encoded as specified for each metric type. For metric types that are smaller than 4 octets in size, the most significant bits are filled with zeros. The percentage metric margin is encoded as an unsigned integer percentage value.

#### 5.7.3. SR Segment List Bandwidth Sub-TLV

The SR Segment List Bandwidth sub-TLV is an optional sub-TLV used to report the bandwidth allocated to the specific SID-List by the path computation entity. Only a single instance of this sub-TLV is advertised for a given Segment List. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

It is a sub-TLV of the SR Segment List TLV and has the following format:

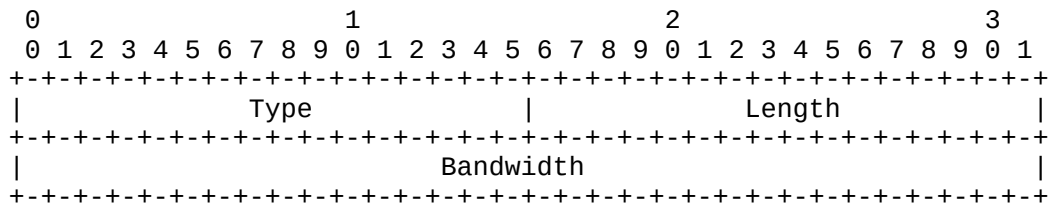


Figure 32 SR Segment List Bandwidth Sub-TLV Format

Where:

* Type: 1216

* Length: 4 octets

* Bandwidth: 4 octets which specify the allocated bandwidth in unit of bytes per second in IEEE floating point format [IEEE754].

### 5.7.4. SR Segment List Identifier Sub-TLV

The SR Segment List Identifier sub-TLV is an optional sub-TLV used to report an identifier associated with the specific SID-List. Only a single instance of this sub-TLV is advertised for a given Segment List. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

It is a sub-TLV of the SR Segment List TLV and has the following format:

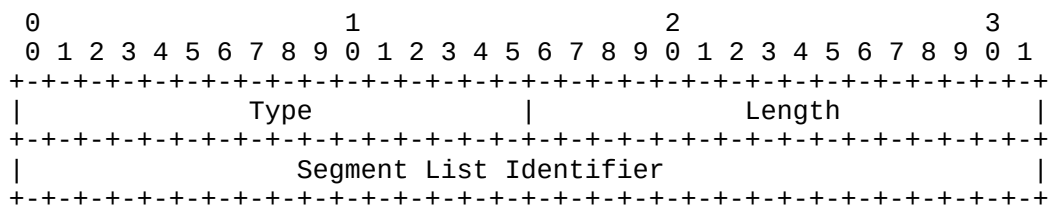


Figure 33 SR Segment List Identifier Sub-TLV Format

Where:

* Type: 1217

* Length: 4 octets

* Segment List Identifier: 4 octets which carry a 32-bit unsigned non-zero integer that serves as the identifier associated with the segment list. A value of 0 indicates that there is no identifier associated with the Segment List. The scope of this identifier is the SR Policy Candidate path.

## 6. Procedures

The BGP-LS advertisements for the SR Policy Candidate Path NLRI type are generally originated by the headend node for the SR Policies that are instantiated on its local node (i.e., the headend is the BGP-LS Producer). The BGP-LS Producer may also be a node (e.g., a PCE) that is advertising on behalf of the headend.

For the reporting of SR Policy Candidate Paths, the NLRI descriptor TLV as specified in Section 4 is used. An SR Policy candidate path may be instantiated on the headend node via a local configuration, PCEP, or BGP SR Policy signaling and this is indicated via the SR Protocol Origin. When a PCE node is the BGP-LS Producer, it uses the "in PCEP" variants of the SR Protocol Origin (where available) so as to distinguish them from advertisements by headend nodes. The SR Policy Candidate Path’s state and attributes are encoded in the BGP-LS Attribute field as SR Policy State TLVs and sub-TLVs as described in Section 5. The SR Candidate Path State TLV as defined in Section 5.3 is included to report the state of the candidate path. The SR BSID TLV as defined in Section 5.1 or Section 5.2 is included to report the BSID of the candidate path when one is either specified or allocated by the headend. The constraints and the optimization metric for the SR Policy Candidate Path are reported using the SR Candidate Path Constraints TLV and its sub-TLVs as described in Section 5.6. The SR Segment List TLV is included for each of the SID-List(s) associated with the candidate path. Each SR Segment List TLV in turn includes SR Segment sub-TLV(s) to report the segment(s) and their status. The SR Segment List Metric sub-TLV is used to report the metric values at an individual SID List level.

## 7. Manageability Considerations

The Existing BGP operational and management procedures apply to this document. No new procedures are defined in this document. The considerations as specified in [RFC9552] apply to this document.

In general, the SR Policy head-end nodes are responsible for the advertisement of SR Policy state information.

## 8. IANA Considerations

This section describes the code point allocation by IANA for this document.

### 8.1. BGP-LS NLRI-Types

IANA maintains a registry called "BGP-LS NLRI-Types" in the "Border Gateway Protocol - Link State (BGP-LS) Parameters" registry group.

The following table lists the code points that have been allocated by IANA:

| Type | NLRI Type | Reference |
|---|---|---|
| 5 | SR Policy Candidate Path NLRI | this document |

Table 2 NLRI Type Codepoint

### 8.2. BGP-LS Protocol-IDs

IANA maintains a registry called "BGP-LS Protocol-IDs" in the "Border Gateway Protocol - Link State (BGP-LS) Parameters" registry group.

The following Protocol-ID codepoints have been allocated by IANA:

| Protocol-ID | NLRI information source protocol | Reference |
|---|---|---|
| 9 | Segment Routing | this document |

Table 3 Protocol ID Codepoint

### 8.3. BGP-LS TLVs

IANA maintains a registry called "BGP-LS NLRI and Attribute TLVs" in the "Border Gateway Protocol - Link State (BGP-LS) Parameters" registry group.

The following table lists the TLV code points that have been allocated by IANA:

| Code Point | Description | Value defined in |
|---|---|---|
| 554 | SR Policy Candidate Path Descriptor | this document |
| 1201 | SR Binding SID | this document |
| 1202 | SR Candidate Path State | this document |
| 1203 | SR Candidate Path Name | this document |
| 1204 | SR Candidate Path Constraints | this document |
| 1205 | SR Segment List | this document |
| 1206 | SR Segment | this document |
| 1207 | SR Segment List Metric | this document |
| 1208 | SR Affinity Constraint | this document |
| 1209 | SR SRLG Constraint | this document |
| 1210 | SR Bandwidth Constraint | this document |
| 1211 | SR Disjoint Group Constraint | this document |
| 1212 | SRv6 Binding SID | this document |
| 1213 | SR Policy Name | this document |
| 1214 | SR Bidirectional Group Constraint | this document |
| 1215 | SR Metric Constraint | this document |
| 1216 | SR Segment List Bandwidth | this document |
| 1217 | SR Segment List Identifier | this document |

Table 4  NLRI and Attribute TLVs Codepoint

## 8.4. SR Policy Protocol Origin

Note to IANA (RFC editor to remove this before publication): The new registry creation request below is also present in the draft-ietf-pce-segment-routing-policy-cp.  IANA is requested to process the registry creation via the first of these two documents to reach publication stage and the authors of the other document would update the IANA considerations suitably.  The initial allocations in this document are a super-set of the initial allocations in draft-ietf-pce-segment-routing-policy-cp.

This document requests IANA to maintain a new registry under "Segment Routing" registry group with the allocation policy of "Expert Review" [RFC8126] using the guidelines for Designated Experts as specified in [RFC9256].  The new registry is called "SR Policy Protocol Origin" and should have the reference to this document.  This registry contains the codepoints allocated to the "Protocol Origin" field defined in Section 4.

The registry contains the following codepoints, with initial values, to be assigned by IANA with the reference set to this document:

| Code Point | Protocol Origin | Reference |
|---|---|---|
| 0 | Reserved (not to be used) | this document |
| 1 | PCEP | this document |
| 2 | BGP SR Policy | this document |
| 3 | Configuration (CLI, YANG model via NETCONF, etc.) | this document |
| 4-9 | Unassigned | this document |
| 10 | PCEP (In PCEP or when BGP-LS Producer is PCE) | this document |
| 11-19 | Unassigned | this document |
| 20 | BGP SR Policy (In PCEP or when BGP-LS Producer is PCE) | this document |
| 21-29 | Unassigned | this document |
| 30 | Configuration (CLI, YANG model via NETCONF, etc.) (In PCEP or when BGP-LS Producer is PCE) | this document |
| 31-250 | Unassigned | this document |
| 251-255 | Private Use (not to be assigned by IANA) | this document |

Table 5  SR Policy Protocol Origin Codepoint

## 8.5.  BGP-LS SR Segment Descriptors

This document requests IANA to create a registry called "SR Segment Descriptor Types" under the "Border Gateway Protocol - Link State (BGP-LS) Parameters" registry group with the allocation policy of "Expert Review" [RFC8126] using the guidelines for Designated Experts as specified in [RFC9552].  There is also an additional guideline to the Designated Experts to maintain the alignment between the allocations in this registry with those in the "Segment Types" registry under the "Segment Routing" registry group.  This requires that an allocation in the Segment Routing "Segment Types" registry is required before allocation can be done in the BGP-LS "SR Segment Descriptor Types" registry for a new segment type.  However, this does not mandate that the specification of a new Segment Routing Segment Type also requires the specification of its equivalent SR Segment Descriptor Type in BGP-LS; that can be done as and when required while maintaining alignment.

This registry contains the codepoints allocated to the "Segment Type" field defined in Section 5.7.1 and described in Section 5.7.1.1.  The registry contains the following codepoints, with initial values, to be assigned by IANA with the reference set to this document:

| Code Point | Segment Description | Reference |
|---|---|---|
| 0 | Reserved (not to be used) | this document |
| 1 | (Type A) SR-MPLS Label | this document |
| 2 | (Type B) SRv6 SID as IPv6 address | this document |
| 3 | (Type C) SR-MPLS Prefix SID as IPv4 Node Address | this document |
| 4 | (Type D) SR-MPLS Prefix SID as IPv6 Node Global Address | this document |
| 5 | (Type E) SR-MPLS Adjacency SID as IPv4 Node Address & Local Interface ID | this document |
| 6 | (Type F) SR-MPLS Adjacency SID as IPv4 Local & Remote Interface Addresses | this document |
| 7 | (Type G) SR-MPLS Adjacency SID as pair of IPv6 Global Address & Interface ID for Local & Remote nodes | this document |
| 8 | (Type H) SR-MPLS Adjacency SID as pair of IPv6 Global Addresses for the Local & Remote Interface | this document |
| 9 | (Type I) SRv6 END SID as IPv6 Node Global Address | this document |
| 10 | (Type J) SRv6 END.X SID as pair of IPv6 Global Address & Interface ID for Local & Remote nodes | this document |
| 11 | (Type K) SRv6 END.X SID as pair of IPv6 Global Addresses for the Local & Remote Interface | this document |
| 12-255 | Unassigned | this document |

Table 6  SR Segment Descriptor Types Codepoint

## 8.6.  BGP-LS SR Policy Metric Type

This document requests IANA to create a registry called "BGP-LS SR Policy Metric Type" under the "Border Gateway Protocol - Link State (BGP-LS) Parameters" registry group with the allocation policy of "Expert Review" [RFC8126] using the guidelines for Designated Experts as specified in [RFC9552].  This registry contains the codepoints allocated to the "metric type" field defined in Section 5.7.2.  The registry contains the following codepoints, with initial values, to be assigned by IANA with the reference set to this document:

| Code Point | Metric Type | Reference |
|---|---|---|
| 0 | IGP | this document |
| 1 | Min Unidirectional Delay | this document |
| 2 | TE | this document |
| 3 | Hop Count | this document |
| 4 | SID List Length | this document |
| 5 | Bandwidth | this document |
| 6 | Avg Unidirectional Delay | this document |
| 7 | Unidirectional Delay Variation | this document |
| 8 | Loss | this document |
| 9-127 | Unassigned | this document |
| 128-255 | User Defined | this document |

Table 7 SR Policy Metric Type Codepoint

## 9. Security Considerations

Procedures and protocol extensions defined in this document do not affect the base BGP security model. See [RFC6952] for details. The security considerations of the base BGP-LS specification as described in [RFC9552] also apply.

The BGP-LS 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] and its security considerations also apply to BGP sessions when carrying SR Policy information. The SR Policies advertised to controllers and other applications via BGP-LS are expected to be used entirely within this trusted SR domain, i.e., within a single AS or between 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 and/or controllers/applications in a secure manner within this trusted SR domain. The general guidance for BGP-LS with respect to isolation of BGP-LS sessions from BGP sessions for other address-families (refer security considerations of [RFC9552]) may be used to ensure that the SR Policy information is not advertised by accident or error to an EBGP peering session outside the SR domain.

Additionally, it may be considered 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.

## 10. Contributors

The following people have substantially contributed to the editing of this document:

Clarence Filsfils
Cisco Systems
Email: cfilsfil@cisco.com

Mach (Guoyi) Chen
Huawei Technologies
Email: mach.chen@huawei.com

## 11. Acknowledgements

The authors would like to thank Dhruv Dhody, Mohammed Abdul Aziz Khalid, Lou Berger, Acee Lindem, Siva Sivabalan, Arjun Sreekantiah, Dhanendra Jain, Francois Clad, Zafar Ali, Stephane Litkowski, Aravind Babu Mahendra Babu, Geetanjalli Bhalla, Ahmed Bashandy, Mike Koldychev, Samuel Sidor, Alex Tokar, Rajesh Melarcode Venkatesswaran, Lin Changwang, Liu Yao, Joel Halpern, and Ned Smith for their review and valuable comments. The authors would also like to thank Susan Hares for her shepherd review of the document and helpful comments to improve this document. The authors would like to thank John Scudder for his AD review and helpful suggestions to improve this document.

## 12. References

### 12.1. Normative References

[I-D.ietf-lsr-flex-algo-bw-con]
Hegde, S., Britto, W., Shetty, R., Decraene, B., Psenak, P., and T. Li, "IGP Flexible Algorithms: Bandwidth, Delay, Metrics and Constraints", Work in Progress, Internet-Draft, draft-ietf-lsr-flex-algo-bw-con-22, 13 February 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-lsr-flex-algo-bw-con-22>.

[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>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <https://www.rfc-editor.org/info/rfc2328>.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <https://www.rfc-editor.org/info/rfc3630>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <https://www.rfc-editor.org/info/rfc5305>.

[RFC5329] Ishiguro, K., Manral, V., Davey, A., and A. Lindem, Ed., "Traffic Engineering Extensions to OSPF Version 3", RFC 5329, DOI 10.17487/RFC5329, September 2008, <https://www.rfc-editor.org/info/rfc5329>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <https://www.rfc-editor.org/info/rfc5440>.

[RFC7471] Giacalone, S., Ward, D., Drake, J., Atlas, A., and S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", RFC 7471, DOI 10.17487/RFC7471, March 2015, <https://www.rfc-editor.org/info/rfc7471>.

[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>.

[RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 8570, DOI 10.17487/RFC8570, March 2019, <https://www.rfc-editor.org/info/rfc8570>.

[RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <https://www.rfc-editor.org/info/rfc8664>.

[RFC8697] Minei, I., Crabbe, E., Sivabalan, S., Ananthakrishnan, H., Dhody, D., and Y. Tanaka, "Path Computation Element Communication Protocol (PCEP) Extensions for Establishing Relationships between Sets of Label Switched Paths (LSPs)", RFC 8697, DOI 10.17487/RFC8697, January 2020, <https://www.rfc-editor.org/info/rfc8697>.

[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>.

[RFC9086] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Patel, K., Ray, S., and J. Dong, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing BGP Egress Peer Engineering", RFC 9086, DOI 10.17487/RFC9086, August 2021, <https://www.rfc-editor.org/info/rfc9086>.

[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>.

[RFC9514] Dawra, G., Filsfils, C., Talaulikar, K., Ed., Chen, M., Bernier, D., and B. Decraene, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing over IPv6 (SRv6)", RFC 9514, DOI 10.17487/RFC9514, December 2023, <https://www.rfc-editor.org/info/rfc9514>.

[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>.

### 12.2. Informative References

[I-D.ietf-idr-bgp-ls-te-path]
Previdi, S., Talaulikar, K., Dong, J., Gredler, H., and J. Tantsura, "Advertisement of Traffic Engineering Paths using BGP Link-State", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-te-path-02, 11 November 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-te-path-02>.

[I-D.ietf-idr-bgp-sr-segtypes-ext]
Talaulikar, K., Filsfils, C., Previdi, S., Mattes, P., and D. Jain, "Segment Routing Segment Types Extensions for BGP SR Policy", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-sr-segtypes-ext-08, 20 February 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-sr-segtypes-ext-08>.

[I-D.ietf-idr-sr-policy-safi]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-13, 6 February 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-13>.

[IEEE754]
Institute of Electrical and Electronics Engineers, "IEEE Standard for Floating-Point Arithmetic", IEEE 754-2019, DOI 10.1109/ieeestd.2019.8766229, 22 July 2019, <https://ieeexplore.ieee.org/document/8766229>.

[RFC2702]
Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, DOI 10.17487/RFC2702, September 1999, <https://www.rfc-editor.org/info/rfc2702>.

[RFC4202]
Kompella, K., Ed. and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, DOI 10.17487/RFC4202, October 2005, <https://www.rfc-editor.org/info/rfc4202>.

[RFC4655]
Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <https://www.rfc-editor.org/info/rfc4655>.

[RFC5065]
Traina, P., McPherson, D., and J. Scudder, "Autonomous System Confederations for BGP", RFC 5065, DOI 10.17487/RFC5065, August 2007, <https://www.rfc-editor.org/info/rfc5065>.

[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>.

[RFC7308] Osborne, E., "Extended Administrative Groups in MPLS Traffic Engineering (MPLS-TE)", RFC 7308, DOI 10.17487/RFC7308, July 2014, <https://www.rfc-editor.org/info/rfc7308>.

[RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <https://www.rfc-editor.org/info/rfc8231>.

[RFC8800] Litkowski, S., Sivabalan, S., Barth, C., and M. Negi, "Path Computation Element Communication Protocol (PCEP) Extension for Label Switched Path (LSP) Diversity Constraint Signaling", RFC 8800, DOI 10.17487/RFC8800, July 2020, <https://www.rfc-editor.org/info/rfc8800>.

## Authors' Addresses

Stefano Previdi
Individual
Email: stefano@previdi.net

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

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing
100095
China
Email: jie.dong@huawei.com

Hannes Gredler
RtBrick Inc.
Email: hannes@rtbrick.com

Jeff Tantsura
Nvidia
Email: jefftant.ietf@gmail.com