

Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: 24 January 2024

S. Previdi
Huawei Technologies
K. Talaulikar, Ed.
Cisco Systems
J. Dong
Huawei Technologies
H. Gredler
RtBrick Inc.
J. Tantsura
Microsoft
23 July 2023

**Advertisement of Segment Routing Policies using BGP Link-State
draft-ietf-idr-bgp-ls-sr-policy-01**

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 24 January 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

| | | |
|------------------------|---|--------------------|
| 1. | Introduction | 3 |
| 1.1. | Requirements Language | 5 |
| 2. | Carrying SR Policy Information in BGP | 5 |
| 3. | SR Policy Candidate Path NLRI Type | 6 |
| 4. | SR Policy Candidate Path Descriptor | 7 |
| 5. | SR Policy State TLVs | 9 |
| 5.1. | SR Binding SID TLV | 9 |
| 5.2. | SRV6 Binding SID TLV | 11 |
| 5.3. | SR Candidate Path State TLV | 13 |
| 5.4. | SR Policy Name TLV | 15 |
| 5.5. | SR Candidate Path Name TLV | 15 |
| 5.6. | SR Candidate Path Constraints TLV | 16 |
| 5.6.1. | SR Affinity Constraint Sub-TLV | 18 |
| 5.6.2. | SR SRLG Constraint Sub-TLV | 20 |
| 5.6.3. | SR Bandwidth Constraint Sub-TLV | 20 |
| 5.6.4. | SR Disjoint Group Constraint Sub-TLV | 21 |
| 5.6.5. | SR Bidirectional Group Constraint Sub-TLV | 23 |
| 5.6.6. | SR Metric Constraint Sub-TLV | 24 |
| 5.7. | SR Segment List TLV | 26 |
| 5.8. | SR Segment Sub-TLV | 28 |
| 5.8.1. | Segment Descriptors | 30 |
| 5.9. | SR Segment List Metric Sub-TLV | 37 |
| 5.10. | SR Segment List Bandwidth Sub-TLV | 39 |
| 5.11. | SR Segment List Identifier Sub-TLV | 39 |
| 6. | Procedures | 40 |
| 7. | Manageability Considerations | 41 |
| 8. | IANA Considerations | 41 |
| 8.1. | BGP-LS NLRI-Types | 41 |
| 8.2. | BGP-LS Protocol-IDs | 42 |
| 8.3. | BGP-LS TLVs | 42 |
| 8.4. | BGP-LS SR Policy Protocol Origin | 43 |
| 8.5. | BGP-LS SR Segment Descriptors | 44 |
| 8.6. | BGP-LS Metric Type | 45 |
| 9. | Security Considerations | 46 |
| 10. | Contributors | 46 |
| 11. | Acknowledgements | 46 |
| 12. | References | 46 |

[12.1](#). Normative References [46](#)
[12.2](#). Informative References [48](#)
 Authors' Addresses [49](#)

1. Introduction

SR Policy architecture details are specified in [[RFC9256](#)]. An SR Policy comprises one or more candidate paths (CP) 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 CP in turn may have one or more SID-List of which one or more may be active; when multiple are active then traffic is load balanced over them. This document covers the advertisement of state information at the individual SR Policy CP 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 a controller which allows the network operator to optimize several functions and operations through the use of a controller aware of both topology and state information.

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 a SR Policy state collection mechanism complementary to the mechanism defined in [[RFC8231](#)].

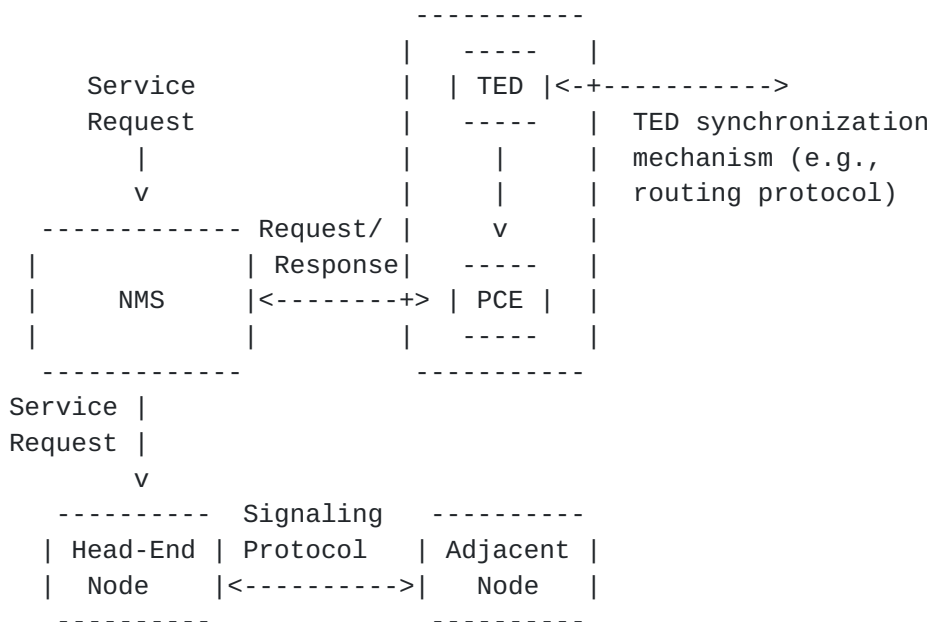


Figure 1. Management-Based PCE Usage

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 SR Policy paths' state 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 [\[I-D.ietf-idr-rfc7752bis\]](#). 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. While this document focuses on SR Policies, #draft-ietf-idr-bgp-ls-te-lsp# introduces further extension to support other TE Paths such as MPLS-TE LSPs.

This extensions specified in this document complement the BGP SR Policy SAFI [I-D.ietf-idr-segment-routing-te-policy] that is 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.


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 [I-D.ietf-idr-rfc7752bis] is extended to carry the SR Policy information. New TLVs carried in the BGP Link-State Attribute defined in [I-D.ietf-idr-rfc7752bis] are also defined to carry the attributes of a SR Policy in the subsequent sections.

The format of "Link-State NLRI" is defined in [I-D.ietf-idr-rfc7752bis] as follows:

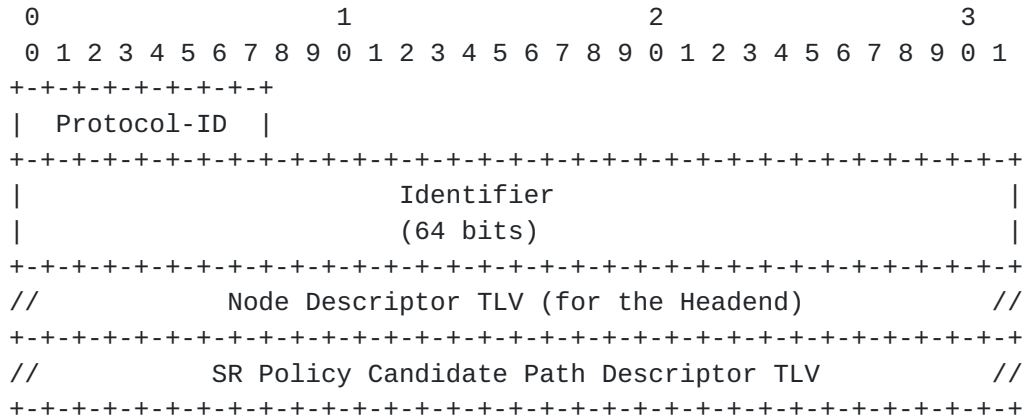


An additional "NLRI Type" known as SR Policy Candidate Path NLRI (value 5) is defined for the advertisement of SR Policy Information.

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:



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 to be used for advertisement of SR Policies.
- * "Identifier" is an 8 octet value as defined in [\[I-D.ietf-idr-rfc7752bis\]](#).
- * "Local Node Descriptor" (TLV 256) [\[I-D.ietf-idr-rfc7752bis\]](#) 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 MUST include the following Node Descriptor TLVs only 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) [\[I-D.ietf-idr-rfc7752bis\]](#), 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 MUST include at least one of the following Node Descriptor TLVs:

- * IPv4 Router-ID of Local Node (TLV 1028) [[I-D.ietf-idr-rfc7752bis](#)], which contains the IPv4 TE Router-ID of the headend node of the SR Policy.
- * IPv6 Router-ID of Local Node (TLV 1029) [[I-D.ietf-idr-rfc7752bis](#)], which contains the IPv6 TE Router-ID of the headend node of the SR Policy.

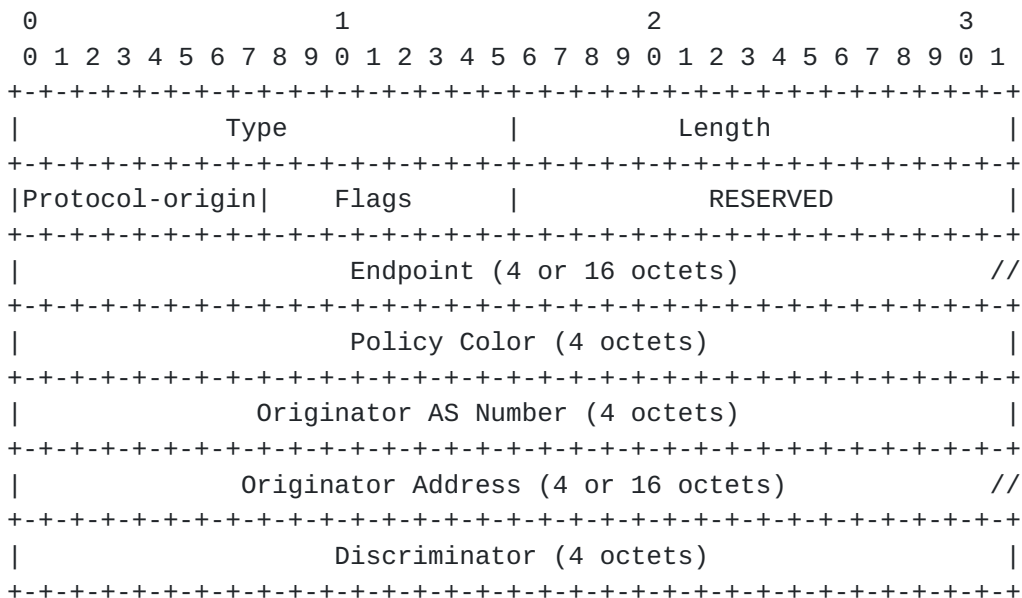
The Local Node Descriptor TLV MAY include the following Node Descriptor TLVs only when the headend node is the BGP-LS Producer:

- * 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 [[I-D.ietf-idr-rfc7752bis](#)] 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 shall always 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. 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.


4. SR Policy Candidate Path Descriptor

The SR Policy Candidate Path Descriptor TLV identifies a Segment Routing Policy Candidate Path (CP) as defined in [[RFC9256](#)]. It is used with the Protocol-ID set to Segment Routing to advertise the SR Policy Candidate Path NLRI Type. It is a mandatory TLV for SR Policy Candidate Path NLRI type. The TLV has the following 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. The protocol-origin codepoints defined in this document 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
- * Color: 4 octets that indicate the color of the SR Policy
- * 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

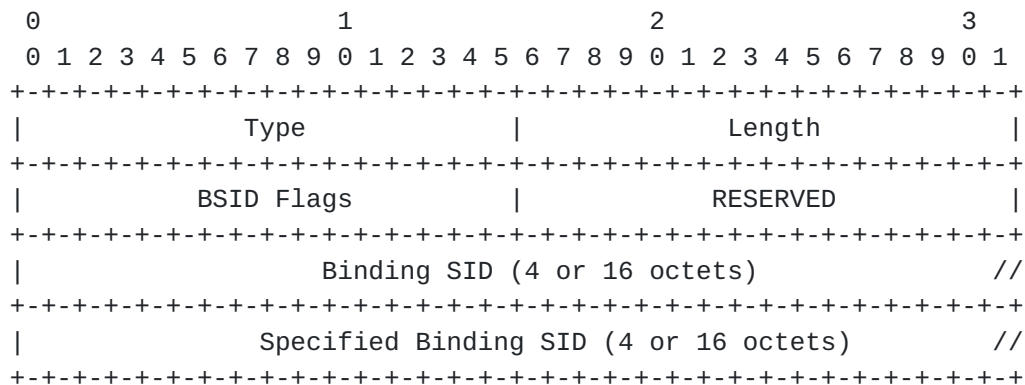
This section defines the various TLVs which enable the headend to report the state at the SR Policy CP level. These TLVs (and their sub-TLVs) are carried in the optional non-transitive BGP-LS Attribute defined in [[I-D.ietf-idr-rfc7752bis](#)] associated with the SR Policy CP NLRI type.

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 CP. 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 CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.

The TLV has the following 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 CP. 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 dataplane for the BSIDs and if they are 16 octet SRv6 SID when set and 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.
- 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.
- * 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 are 4-octet carrying the MPLS Label or 16-octet carrying the SRv6 SID based on the BSID D-flag. 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.

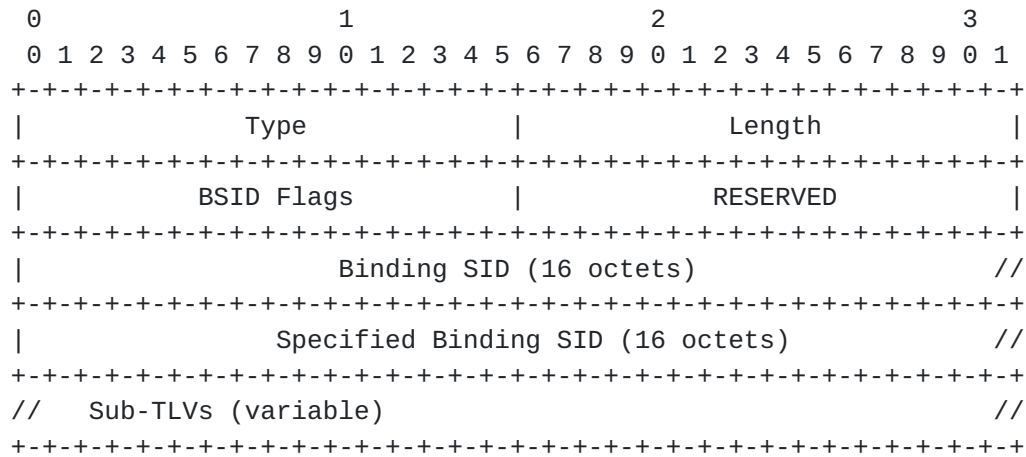


In the case of an SRv6, the SR Binding SID sub-TLV does not have the ability to signal the SRv6 Endpoint Behavior [[RFC8986](#)] or the structure of the SID. Therefore, the SR Binding SID sub-TLV SHOULD NOT be used for the advertisement of an SRv6 Binding SID. Instead, the SRv6 Binding SID TLV defined 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 is supported 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 CP. 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 CP.

The TLV has the following format:



where:

- * Type: 1212
- * Length: variable
- * BSID Flags: 2-octet field that indicates attribute and status of the Binding SID (BSID) associated with this CP. 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:

- 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.
- F-Flag: Indicates the BSID value is one allocated dynamically due to fallback (e.g. when specified BSID is unavailable) when set.

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

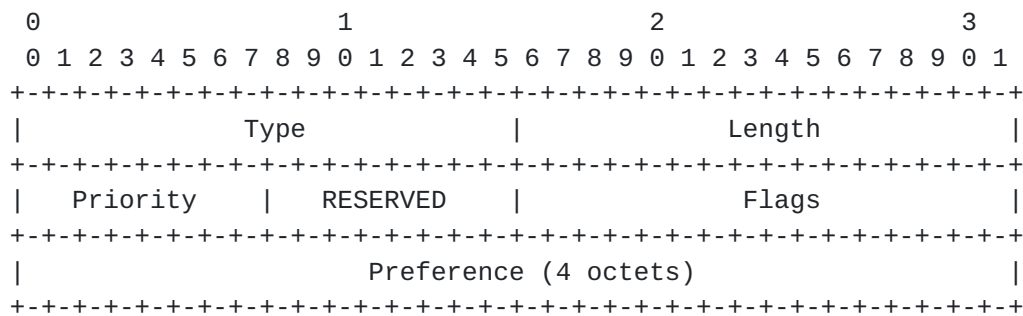
- * Binding SID: It indicates the operational or allocated BSID value based on the status flags.
- * Specified BSID: It is used to report the explicitly specified BSID value regardless of whether it is successfully allocated or not. The field is set to value 0 when BSID has not been specified.
- * Sub-TLVs: 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) defined in [[I-D.ietf-idr-bgppls-srv6-ext](#)] are used as sub-TLVs of the SRv6 Binding SID TLV to optionally indicate the SRv6 Endpoint behavior and SID structure for the Binding SID value in the TLV.

5.3. SR Candidate Path State TLV

The SR Candidate Path (CP) State TLV provides the operational status and attributes of the SR Policy at the CP level. Only a single instance of this TLV is advertised for a given CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.


The TLV has the following format:



where:

- * Type: 1202
- * Length: 8 octets
- * Priority: 1-octet value which indicates the priority of the CP. 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 CP. 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:


- S-Flag: Indicates the CP is in an administrative shut state when set
- A-Flag: Indicates the CP is the active path (i.e. one provisioned in the forwarding plane) for the SR Policy when set
- B-Flag: Indicates the CP is the backup path (i.e. one identified for path protection of the active path) for the SR Policy when set
- E-Flag: Indicates that the CP has been evaluated for validity (e.g. headend may evaluate CPs based on their preferences) when set
- V-Flag: Indicates the CP has at least one valid SID-List when set. When the E-Flag is clear (i.e. the CP has not been evaluated), then this flag MUST be set to 0 by the originator and ignored by the receiver.
- O-Flag: Indicates the CP was instantiated by the headend due to an on-demand nexthop trigger based on a local template when set. Refer to [section 8.5 of \[RFC9256\]](#) for details.
- D-Flag: Indicates the CP was delegated for computation to a PCE/controller when set
- C-Flag: Indicates the CP was provisioned by a PCE/controller when set
- I-Flag: Indicates the CP will perform the "drop upon invalid" behavior when no other active path is available for this SR Policy and this path is the one with the best preference amongst the available CPs. Refer to [section 8.2 of \[RFC9256\]](#) for details.

- T-Flag: Indicates the CP has been marked as eligible for use as Transit Policy on the headend when set. Refer to [section 8.3 of \[RFC9256\]](#).
 - U-Flag: Indicates the SR Policy that the CP belongs to is dropping traffic as a result of the "drop upon invalid" behavior being activated.
- * Preference: 4-octet value which indicates the preference of the CP. 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 CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.

The TLV has the following format:




where:

- * Type: 1213
- * Length: variable
- * SR Policy Name: Symbolic name for the SR Policy 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.

5.5. 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 CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.

The TLV has the following format:



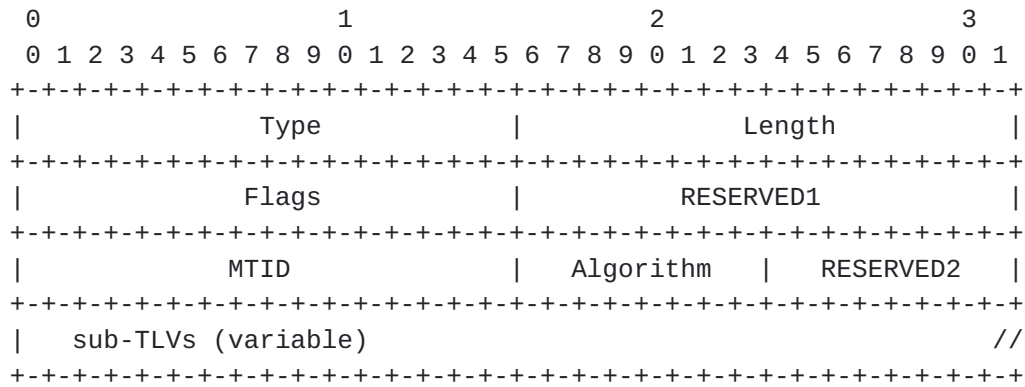
where:

- * 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 CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.

The TLV has the following format:



where:

- * Type: 1204
- * Length: variable
- * Flags: 2-octet field that indicates the constraints that are being applied to the CP. 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 CP uses SRV6 dataplane when set and SR/MPLS dataplane when clear
- P-Flag: Indicates that the CP prefers the use of only protected SIDs when set. This flag is mutually exclusive with the U-Flag.
- U-Flag: Indicates that the CP prefers the use of only unprotected SIDs when set. This flag is mutually exclusive with the P-Flag.
- A-Flag: Indicates that the CP uses only the SIDs belonging to the specified SR Algorithm when set
- T-Flag: Indicates that the CP uses only the SIDs belonging to the specified topology when set

- 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
- F-Flag: Indicates that the CP is fixed once computed and not modified except on operator intervention.
- H-Flag: Indicates that the CP uses only adjacency SIDs and traverses hop-by-hop over the links corresponding to those adjacency SIDs when set
- * 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.
- * RESERVED2: 1 octet. MUST be set to 0 by the originator and MUST be ignored by a receiver.
- * sub-TLVs: optional sub-TLVs MAY be included in this TLV to describe other constraints.

The following constraint sub-TLVs are defined for the SR CP Constraints TLV.

5.6.1. SR Affinity Constraint Sub-TLV

The SR Affinity Constraint sub-TLV is an optional sub-TLV of the SR CP 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 CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.

The sub-TLV has the following format:

| 0 | | | | | | | | | | 1 | | | | | | | | | | 2 | | | | | | | | | | 3 | | | | | | | | | |
|--------------------------------------|---|---|---|---|---|---|---|---|---|--------------------------------------|---|---|---|---|---|---|---|---|---|--------------------------------------|---|---|---|---|---|---|---|---|---|----|---|---|---|---|---|---|---|---|---|
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| Type | | | | | | | | | | Length | | | | | | | | | | RESERVED | | | | | | | | | | | | | | | | | | | |
| Exclude-Any EAG (optional, variable) | | | | | | | | | | Include-Any EAG (optional, variable) | | | | | | | | | | Include-All EAG (optional, variable) | | | | | | | | | | // | | | | | | | | | |


where:

- * Type: 1208
- * Length: variable, dependent on the size of the Extended Admin Group. MUST be a 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 CP 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 CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.

The sub-TLV has the following format:




where:

- * Type: 1209
- * Length: variable, dependent on the number of SRLGs encoded. MUST be a multiple of 4 octets.
- * SRLG Values: One or more SRLG values (each 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 CP 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 CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.

The sub-TLV has the following format:



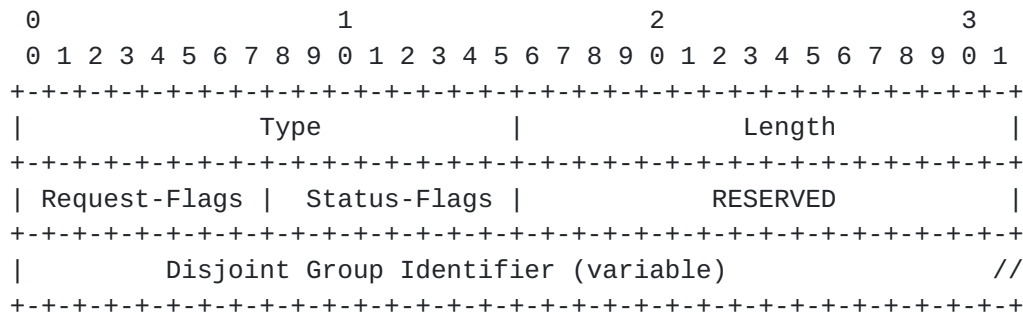
where:

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

5.6.4. SR Disjoint Group Constraint Sub-TLV

The SR Disjoint Group Constraint sub-TLV is an optional sub-TLV of the SR CP 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 computation is expected to achieve the highest level of disjointness requested and when that is not possible then fallback to a lesser level progressively based on the levels indicated. Only a single instance of this sub-TLV is advertised for a given CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.

The sub-TLV has the following 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.


```

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

```

where:

- S-Flag: Indicates that SRLG disjointness is requested when set
- N-Flag: Indicates that node disjointness is requested when set
- L-Flag: Indicates that link disjointness is requested when set
- F-Flag: Indicates that the computation may fallback to a lower level of disjointness amongst the ones requested when all cannot be achieved when set
- I-Flag: Indicates that the computation may fallback to the default best path (e.g. IGP path) in case of none of the desired disjointness can be achieved when set

* 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.

```

  0 1 2 3 4 5 6 7
+--+--+--+--+--+--+
|S|N|L|F|I|X|  |
+--+--+--+--+--+--+

```

where:

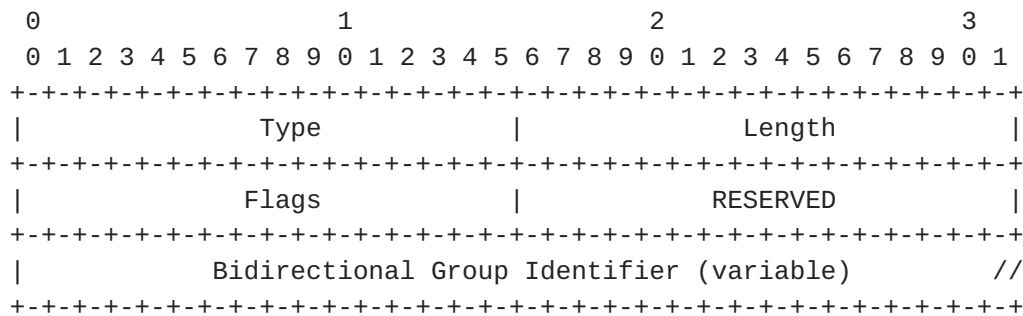
- S-Flag: Indicates that SRLG disjointness is achieved when set
- N-Flag: Indicates that node disjointness is achieved when set
- L-Flag: Indicates that link disjointness is achieved when set
- F-Flag: Indicates that the computation has fallen back to a lower level of disjointness than requested when set
- 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
- X-Flag : Indicates that the disjointness constraint could not be achieved and hence path has been invalidated when set

- * RESERVED: 2 octets. MUST be set to 0 by the originator and MUST be ignored by a receiver.
- * Disjointness Group Identifier: 4-octet value that is the group identifier for a set of disjoint paths. A PCEP Association Object [RFC8697] (including its optional TLVs) MAY also be advertised to convey the disjoint group identifier.

5.6.5. SR Bidirectional Group Constraint Sub-TLV


The SR Bidirectional Group Constraint sub-TLV is an optional sub-TLV of the SR CP 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 CP. If multiple instances are present, then the first one is considered valid and the rest are ignored.

The sub-TLV has the following format:



where:

- * Type: TBD
- * 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 CP of the SR Policy forms the reverse path when set and otherwise it is the forward path when clear
- C-Flag: Indicates that the bidirectional path is co-routed when set
- * 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. A PCEP Association Object [[RFC8697](#)] (including its optional TLVs) MAY also be advertised to convey the bidirectional group identifier.

5.6.6. SR Metric Constraint Sub-TLV


The SR Metric Constraint sub-TLV is an optional sub-TLV of the SR CP Constraints TLV that is used to report the optimization metric of the CP. 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:



where:

- * Type: TBD
- * Length: 12 octets
- * Metric Type: 1-octet field which identifies the type of the metric being used. The metric type code points are listed in [Section 8.6](#) 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:

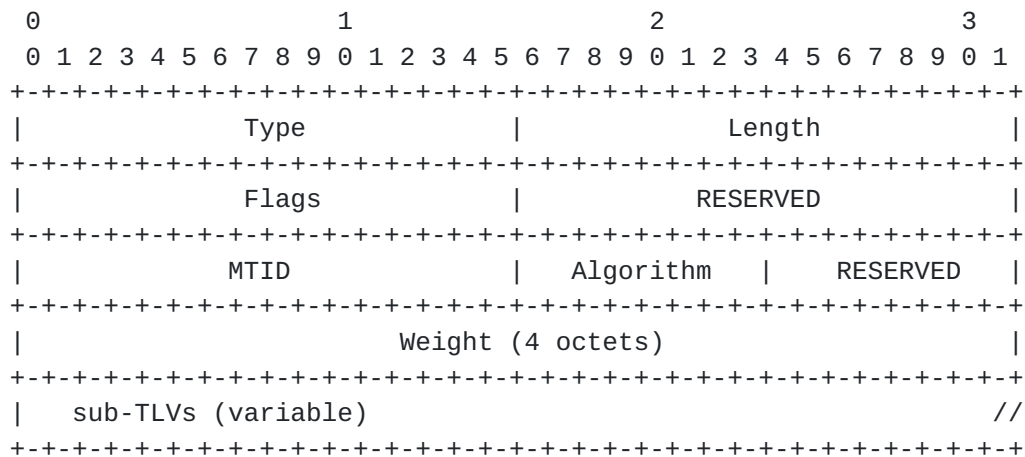
- O-Flag: Indicates that this is the optimization metric being reported for a dynamic CP when set. This bit MUST NOT be set in more than one instance of this TLV for a given CP advertisement.
- M-Flag: Indicates that the metric margin allowed is specified when set.
- A-Flag: Indicates that the metric 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.

- * 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 as either an absolute value or as a percentage of the minimum computed path metric based on the A-flag. The metric margin loosens the criteria for minimum metric path calculation up to the specified metric to accommodate for other factors 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 maximum metric 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.

5.7. SR Segment List TLV

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


The TLV has the following 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:

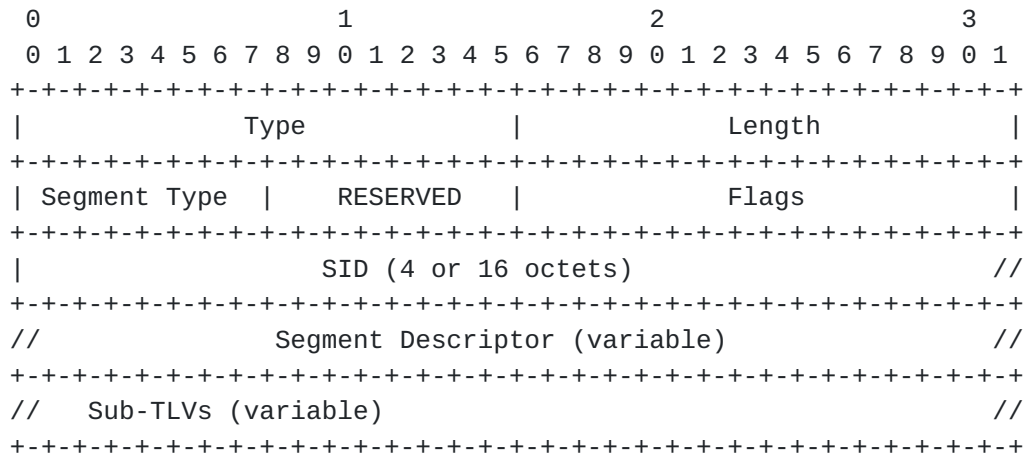
- D-Flag: Indicates the SID-List is comprised of SRv6 SIDs when set and indicates it is comprised of SR/MPLS labels when clear.
- E-Flag: Indicates that SID-List is associated with an explicit candidate path when set and with a dynamic candidate path when clear. All segment lists of a given candidate 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.
- 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.
- T-Flag: Indicates that all the SIDs in the SID-List belong to the specified topology (identified by the multi-topology ID) when set.
- 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.
- * 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 5.8](#)) 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 cases (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, the SR Segment List TLV SHOULD NOT include any SR Segment sub-TLVs.


[5.8](#). 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:



where:

- * Type: 1206
- * Length: variable
- * Segment Type: 1 octet which indicates the type of segment (refer [Section 5.8.1](#) for details)
- * 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 detail in [\[RFC9256\]](#). 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 and failed verification when clear.
- R-Flag: Indicates the SID has been resolved or did not require resolution (e.g. because it is not the first SID) when set and failed resolution when clear.
- 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 5.8.1](#) 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 [[I-D.ietf-idr-bgppls-srv6-ext](#)] are used as sub-sub-TLVs of the SR Segment sub-TLV to optionally indicate the SRv6 Endpoint behavior and SID structure when advertising the SRv6 specific segment types.

5.8.1. 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 describes previously in [Section 5.8](#).

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 |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

5.8.1.1. Type 1: SR-MPLS Label

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




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.

5.8.1.2. Type 2: SRv6 SID

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




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.

5.8.1.3. Type 3: SR-MPLS Prefix SID for IPv4

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:

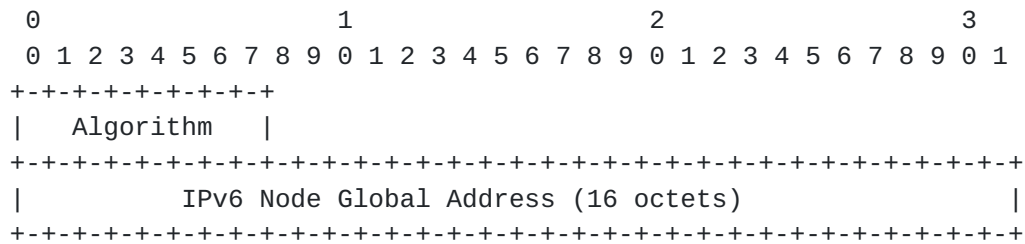


Where:

- * Algorithm: 1-octet value that indicates the algorithm used for picking the SID
- * IPv4 Node Address: 4-octet value which carries the IPv4 address associated with the node

5.8.1.4. Type 4: SR-MPLS Prefix SID for IPv6

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:

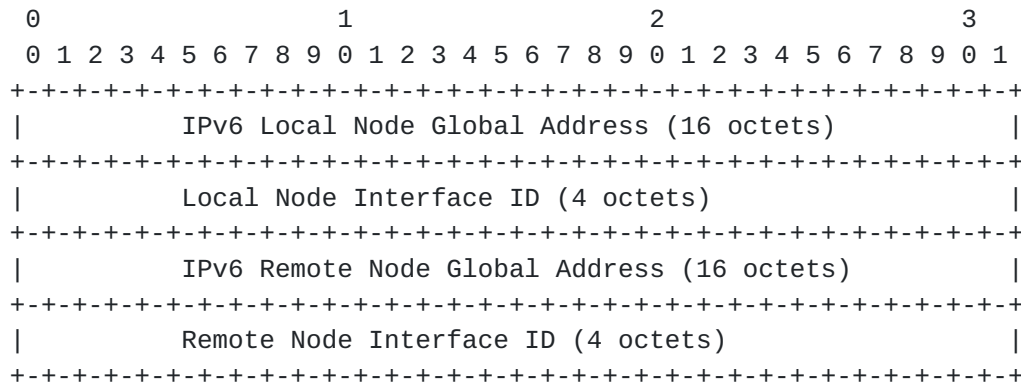


Where:

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

5.8.1.7. Type 7: SR-MPLS Adjacency SID for IPv6 with an interface ID

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:




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).

5.8.1.8. Type 8: SR-MPLS Adjacency SID for IPv6 with an Interface Address

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:

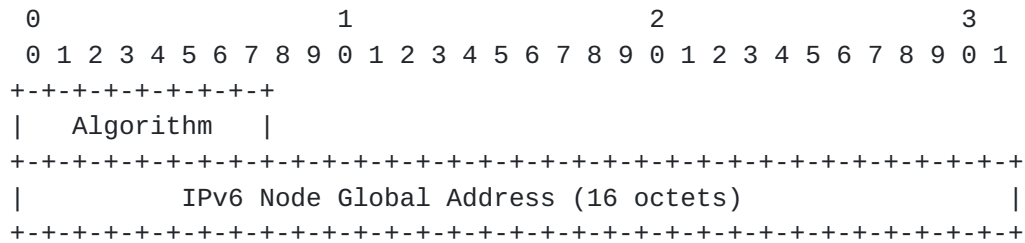


Where:

- * IPv6 Local Address: 16-octet value which carries the local IPv6 address associated with the node
- * IPv6 Remote Address: 16-octet value which carries the remote IPv6 address associated with the node's neighbor

5.8.1.9. Type 9: SRv6 END SID as IPv6 Node Address

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

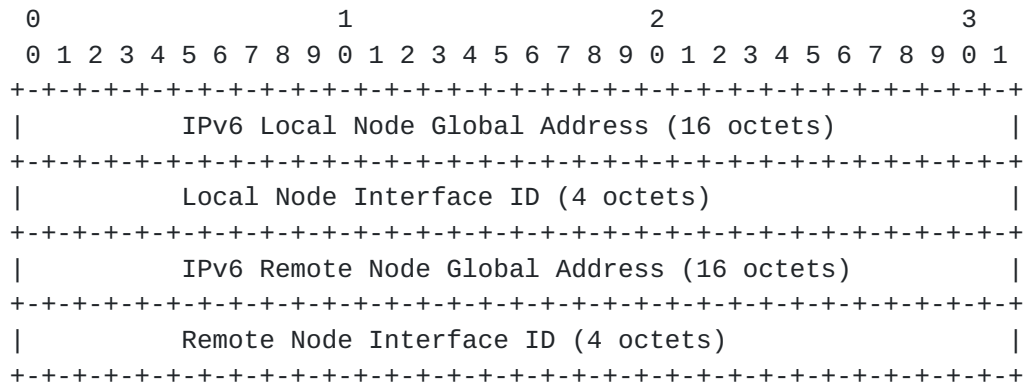


Where:

- * Algorithm: 1-octet value that indicates the algorithm used for picking the SID
- * IPv6 Node Global Address: 16-octet value which carries the IPv6 global address associated with the node

5.8.1.10. Type 10: SRv6 END.X SID as an Anterface ID

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:




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).

5.8.1.11. Type 11: SRv6 END.X SID as an Interface Address

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:



Where:


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

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

5.9. 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 CP level.

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



where:

- * Type: 1207
- * Length: 16 octets
- * Metric Type: 1-octet field which identifies the type of metric. The metric type code points are listed in [Section 8.6](#) 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
 - 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.
 - V-Flag: Indicates that the metric value computed is being reported when set.
- * 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 as either an absolute value or as a percentage of the minimum computed path metric based on the A-flag. The metric margin loosens the criteria for minimum metric path calculation up to the specified metric to accomodate for other factors 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 maximum metric value 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.

5.10. 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 one is considered valid and the rest are ignored.

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




where:

- * Type: TBD
- * Length: 4 octets
- * Bandwidth: 4 octets which specify the allocated bandwidth in unit of bytes per second in IEEE floating point format.

5.11. 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 one is considered valid and the rest are ignored.

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



where:

- * Type: TBD
- * Length: 4 octets
- * Segment List Identifier: 4 octets which carry a 32-bit unsigned non-zero number 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 CP 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 (CP) 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 "reported via PCE" variants of the SR Protocol Origin so as to distinguish them from advertisements by headend nodes. Then 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 CP. The SR BSID TLV as defined in [Section 5.1](#) or [Section 5.2](#) is included to report the BSID of the CP 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 CP. 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 [[I-D.ietf-idr-rfc7752bis](#)] 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 status of code points that have been allocated by IANA:

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

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 |

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 status of TLV code points that have been allocated by IANA and others that are pending allocation:

| Code Point | Description | Value defined in |
|------------|-----------------------------------|------------------|
| 554 | SR Policy CP Descriptor | this document |
| 1201 | SR Binding SID | this document |
| 1202 | SR CP State | this document |
| 1203 | SR CP Name | this document |
| 1204 | SR CP 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 |
| TBD | SR Bidirectional Group Constraint | this document |
| TBD | SR Metric Constraint | this document |
| TBD | SR Segment List Bandwidth | this document |
| TBD | SR Segment List Identifier | this document |

8.4. BGP-LS SR Policy Protocol Origin

This document requests IANA to maintain a new registry under "Segment Routing Parameters" registry group with the allocation policy of "Specification Required" [RFC8126] using the guidelines for Designated Experts as specified in [RFC9256]. The new registry is called "SR Policy Protocol Origin" and 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 |
|------------|--|
| 0 | Reserved (not to be used) |
| 1 | PCEP |
| 2 | BGP SR Policy |
| 3 | Configuration (CLI, YANG model via NETCONF, etc.) |
| 4-9 | Unassigned |
| 10 | PCEP reported via PCE |
| 11-19 | Unassigned |
| 20 | BGP SR Policy reported via PCE |
| 21-29 | Unassigned |
| 30 | Configuration (CLI, YANG model via NETCONF, etc.) reported via PCE |
| 31-250 | Unassigned |
| 251-255 | Private Use (not to be assigned by IANA) |

8.5. BGP-LS SR Segment Descriptors

This document requests IANA to maintain a new registry under "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 [RFC9029]. The new registry is called "SR Segment Descriptor Types" and contains the codepoints allocated to the "Segment Type" field defined in Section 5.8 and described in Section 5.8.1. 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 |
|------------|---|
| 0 | Reserved (not to be used) |
| 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 |
| 12-255 | Unassigned |

8.6. BGP-LS Metric Type

This document requests IANA to maintain a new registry under "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 [RFC9029]. The new registry is called "Metric Type" and contains the codepoints allocated to the "metric type" field defined in Section 5.9. 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 |
|------------|--|
| 0 | IGP Metric |
| 1 | Min Unidirectional Link Delay [RFC7471] |
| 2 | TE Metric [RFC3630] |
| 3 | Hop Count (refer [RFC5440]) |
| 4 | SID List Length |
| 5-250 | Unassigned |
| 251-255 | Private Use (not to be assigned by IANA) |

9. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See [[RFC6952](#)] for details.

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, and Mike Koldychev for their review and valuable comments.

12. References

12.1. Normative References

[I-D.ietf-idr-bgpls-srv6-ext]

Dawra, G., Filsfils, C., Talaulikar, K., Chen, M., Bernier, D., and B. Decraene, "BGP Link State Extensions for SRv6", Work in Progress, Internet-Draft, [draft-ietf-idr-bgpls-srv6-ext-14](#), 17 February 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgpls-srv6-ext-14>>.

[I-D.ietf-idr-rfc7752bis]

Talaulikar, K., "Distribution of Link-State and Traffic Engineering Information Using BGP", Work in Progress, Internet-Draft, [draft-ietf-idr-rfc7752bis-16](#), 20 February 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-rfc7752bis-16>>.

- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC9029] Farrel, A., "Updates to the Allocation Policy for the Border Gateway Protocol - Link State (BGP-LS) Parameters Registries", [RFC 9029](#), DOI 10.17487/RFC9029, June 2021, <<https://www.rfc-editor.org/info/rfc9029>>.
- [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>>.

12.2. Informative References

- [I-D.ietf-idr-segment-routing-te-policy] Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, [draft-ietf-idr-segment-routing-te-policy-20](#), 27 July 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-segment-routing-te-policy-20>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.

Authors' Addresses

Stefano Previdi
Huawei Technologies
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
Microsoft
Email: jefftant.ietf@gmail.com

