

Internet Engineering Task Force (IETF)
Request for Comments: 9252
Category: Standards Track
ISSN: 2070-1721

G. Dawra, Ed.
LinkedIn
K. Talaulikar, Ed.
Cisco Systems
R. Raszuk
NTT Network Innovations
B. Decraene
Orange
S. Zhuang
Huawei Technologies
J. Rabadan
Nokia
July 2022

BGP Overlay Services Based on Segment Routing over IPv6 (SRv6)

Abstract

This document defines procedures and messages for SRv6-based BGP services, including Layer 3 Virtual Private Network (L3VPN), Ethernet VPN (EVPN), and Internet services. It builds on "BGP/MPLS IP Virtual Private Networks (VPNs)" (RFC 4364) and "BGP MPLS-Based Ethernet VPN" (RFC 7432).

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9252>.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction
 - 1.1. Requirements Language
2. SRv6 Services TLVs
3. SRv6 Service Sub-TLVs
 - 3.1. SRv6 SID Information Sub-TLV
 - 3.2. SRv6 Service Data Sub-Sub-TLVs

- 3.2.1. SRv6 SID Structure Sub-Sub-TLV
- 4. Encoding SRv6 SID Information
- 5. BGP-Based L3 Service over SRv6
 - 5.1. IPv4 VPN over SRv6 Core
 - 5.2. IPv6 VPN over SRv6 Core
 - 5.3. Global IPv4 over SRv6 Core
 - 5.4. Global IPv6 over SRv6 Core
- 6. BGP-Based Ethernet VPN (EVPN) over SRv6
 - 6.1. Ethernet Auto-Discovery Route over SRv6 Core
 - 6.1.1. Ethernet A-D per ES Route
 - 6.1.2. Ethernet A-D per EVI Route
 - 6.2. MAC/IP Advertisement Route over SRv6 Core
 - 6.2.1. MAC/IP Advertisement Route with MAC Only
 - 6.2.2. MAC/IP Advertisement Route with MAC+IP
 - 6.3. Inclusive Multicast Ethernet Tag Route over SRv6 Core
 - 6.4. Ethernet Segment Route over SRv6 Core
 - 6.5. IP Prefix Route over SRv6 Core
 - 6.6. EVPN Multicast Routes (Route Types 6, 7, and 8) over SRv6 Core
- 7. Error Handling
- 8. IANA Considerations
 - 8.1. BGP Prefix-SID TLV Types Registry
 - 8.2. SRv6 Service Sub-TLV Types Registry
 - 8.3. SRv6 Service Data Sub-Sub-TLV Types Registry
 - 8.4. BGP SRv6 Service SID Flags Registry
 - 8.5. SAFI Values Registry
- 9. Security Considerations
 - 9.1. Considerations Related to BGP Sessions
 - 9.2. Considerations Related to BGP Services
 - 9.3. Considerations Related to SR over IPv6 Data Plane
- 10. References
 - 10.1. Normative References
 - 10.2. Informative References
- Acknowledgements
- Contributors
- Authors' Addresses

1. Introduction

SRv6 refers to Segment Routing instantiated on the IPv6 data plane [RFC8402].

BGP is used to advertise the reachability of prefixes of a particular service from an egress Provider Edge (PE) to ingress PE nodes.

SRv6-based BGP services refer to the Layer 3 (L3) and Layer 2 (L2) overlay services with BGP as the control plane and SRv6 as the data plane. This document defines procedures and messages for SRv6-based BGP services, including L3VPN, EVPN, and Internet services. It builds on "BGP/MPLS IP Virtual Private Networks (VPNs)" [RFC4364] and "BGP MPLS-Based Ethernet VPN" [RFC7432].

SRv6 SID refers to an SRv6 Segment Identifier, as defined in [RFC8402].

SRv6 Service SID refers to an SRv6 SID associated with one of the service-specific SRv6 Endpoint Behaviors on the advertising PE router, such as (but not limited to) End.DT (look up in the Virtual Routing and Forwarding (VRF) table) or End.DX (cross-connect to a next hop) behaviors in the case of L3VPN service, as defined in [RFC8986]. This document describes how existing BGP messages between PEs may carry SRv6 Service SIDs to interconnect PEs and form VPNs.

To provide SRv6 service with best-effort connectivity, the egress PE signals an SRv6 Service SID with the BGP overlay service route. The ingress PE encapsulates the payload in an outer IPv6 header where the

destination address is the SRv6 Service SID provided by the egress PE. The underlay between the PEs only needs to support plain IPv6 forwarding [RFC8200].

To provide SRv6 service in conjunction with an underlay Service Level Agreement (SLA) from the ingress PE to the egress PE, the egress PE colors the overlay service route with a Color Extended Community [RFC9012] for steering flows for those routes, as specified in Section 8 of [SEGMENT-ROUTING-POLICY]. The ingress PE encapsulates the payload packet in an outer IPv6 header with the SR Policy segment list associated with the related SLA along with the SRv6 Service SID associated with the route using the Segment Routing Header (SRH) [RFC8754]. The underlay nodes whose SRv6 SIDs are part of the SRH segment list MUST support the SRv6 data plane.

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. SRv6 Services TLVs

This document extends the use of the BGP Prefix-SID attribute [RFC8669] to carry SRv6 SIDs and their associated information with the BGP address families that are listed further in this section.

The SRv6 Service TLVs are defined as two new TLVs of the BGP Prefix-SID attribute to achieve signaling of SRv6 SIDs for L3 and L2 services.

SRv6 L3 Service TLV:

This TLV encodes Service SID information for SRv6-based L3 services. It corresponds to the equivalent functionality provided by an MPLS label when received with a Layer 3 service route, as defined in [RFC4364], [RFC4659], [RFC8950], and [RFC9136]. Some SRv6 Endpoint Behaviors that may be encoded are, but not limited to, End.DX4, End.DT4, End.DX6, End.DT6, and End.DT46.

SRv6 L2 Service TLV:

This TLV encodes Service SID information for SRv6-based L2 services. It corresponds to the equivalent functionality provided by an MPLS label for Ethernet VPN (EVPN) Route Types for Layer 2 services, as defined in [RFC7432]. Some SRv6 Endpoint Behaviors that may be encoded are, but not limited to, End.DX2, End.DX2V, End.DT2U, and End.DT2M.

When an egress PE is enabled for BGP Services over the SRv6 data plane, it signals one or more SRv6 Service SIDs enclosed in an SRv6 Service TLV(s) within the BGP Prefix-SID attribute attached to Multiprotocol BGP (MP-BGP) Network Layer Reachability Information (NLRI) defined in [RFC4760], [RFC4659], [RFC8950], [RFC7432], [RFC4364], and [RFC9136], where applicable, as described in Sections 5 and 6.

The support for BGP Multicast VPN (MVPN) Services [RFC6513] with SRv6 is outside the scope of this document.

The following depicts the SRv6 Service TLVs encoded in the BGP Prefix-SID attribute:

[illegible]

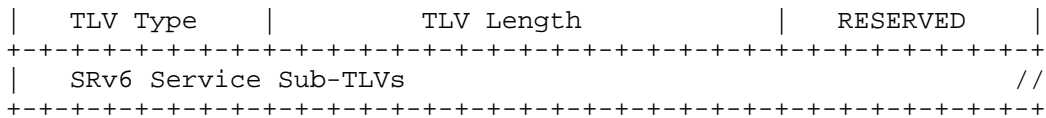


Figure 1: SRv6 Service TLVs

TLV Type (1 octet):
 This field is assigned a value from IANA's "BGP Prefix-SID TLV Types" subregistry. It is set to 5 for the SRv6 L3 Service TLV. It is set to 6 for the SRv6 L2 Service TLV.

TLV Length (2 octets):
 This field specifies the total length, in octets, of the TLV Value.

RESERVED (1 octet):
 This field is reserved; it MUST be set to 0 by the sender and ignored by the receiver.

SRv6 Service Sub-TLVs (variable):
 This field contains SRv6 service-related information and is encoded as an unordered list of Sub-TLVs whose format is described below.

A BGP speaker receiving a route containing the BGP Prefix-SID attribute with one or more SRv6 Service TLVs observes the following rules when advertising the received route to other peers:

- * If the BGP next hop is unchanged during the advertisement, the SRv6 Service TLVs, including any unrecognized Types of Sub-TLV and Sub-Sub-TLV, SHOULD be propagated further. In addition, all Reserved fields in the TLV, Sub-TLV, or Sub-Sub-TLV MUST be propagated unchanged.
- * If the BGP next hop is changed, the TLVs, Sub-TLVs, and Sub-Sub-TLVs SHOULD be updated with the locally allocated SRv6 SID information. Any received Sub-TLVs and Sub-Sub-TLVs that are unrecognized MUST be removed.

3. SRv6 Service Sub-TLVs

The format of a single SRv6 Service Sub-TLV is depicted below:

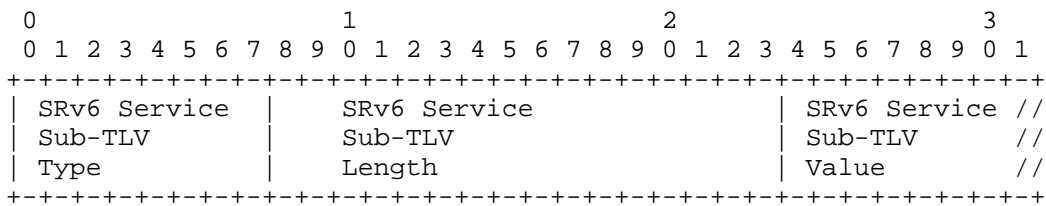


Figure 2: SRv6 Service Sub-TLVs

SRv6 Service Sub-TLV Type (1 octet):
 This field identifies the type of SRv6 service information. It is assigned a value from IANA's "SRv6 Service Sub-TLV Types" subregistry.

SRv6 Service Sub-TLV Length (2 octets):
 This field specifies the total length, in octets, of the Sub-TLV Value field.

SRv6 Service Sub-TLV Value (variable):
 This field contains data specific to the Sub-TLV Type. In addition to fixed-length data, it contains other properties of the

SRv6 service encoded as a set of SRv6 Service Data Sub-Sub-TLVs whose format is described in Section 3.2 below.

3.1. SRv6 SID Information Sub-TLV

SRv6 Service Sub-TLV Type 1 is assigned for the SRv6 SID Information Sub-TLV. This Sub-TLV contains a single SRv6 SID along with its properties. Its encoding is depicted below:

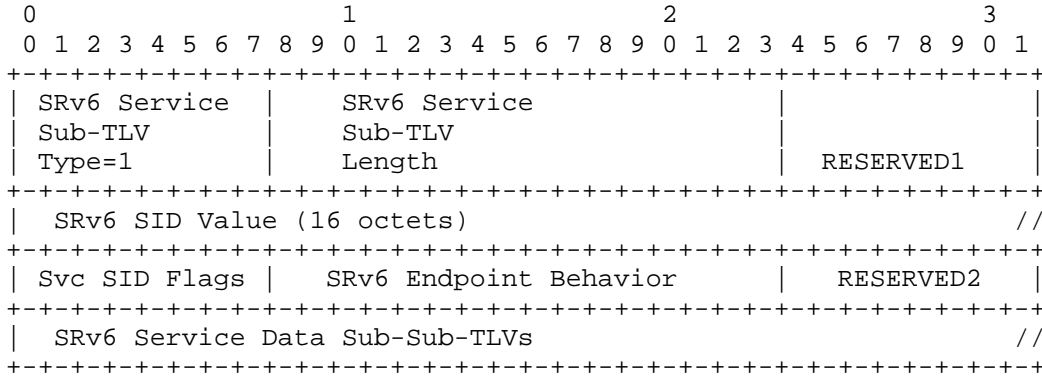


Figure 3: SRv6 SID Information Sub-TLV

SRv6 Service Sub-TLV Type (1 octet):

This field is set to 1 to represent the SRv6 SID Information Sub-TLV.

SRv6 Service Sub-TLV Length (2 octets):

This field contains the total length, in octets, of the Value field of the Sub-TLV.

RESERVED1 (1 octet):

This field MUST be set to 0 by the sender and ignored by the receiver.

SRv6 SID Value (16 octets):

This field encodes an SRv6 SID, as defined in [RFC8986].

SRv6 Service SID Flags (1 octet):

This field encodes SRv6 Service SID Flags -- none are currently defined. It MUST be set to 0 by the sender and any unknown flags MUST be ignored by the receiver.

SRv6 Endpoint Behavior (2 octets):

This field encodes the SRv6 Endpoint Behavior codepoint value that is associated with the SRv6 SID. The codepoints used are from IANA's "SRv6 Endpoint Behaviors" subregistry under the "Segment Routing" registry that was introduced by [RFC8986]. The opaque SRv6 Endpoint Behavior (i.e., value 0xFFFF) MAY be used when the advertising router wishes to abstract the actual behavior of its locally instantiated SRv6 SID.

RESERVED2 (1 octet):

This field MUST be set to 0 by the sender and ignored by the receiver.

SRv6 Service Data Sub-Sub-TLV Value (variable):

This field is used to advertise properties of the SRv6 SID. It is encoded as a set of SRv6 Service Data Sub-Sub-TLVs.

The choice of SRv6 Endpoint Behavior of the SRv6 SID is entirely up to the originator of the advertisement. While Sections 5 and 6 list the SRv6 Endpoint Behaviors that are normally expected to be used by the specific route advertisements, the reception of other SRv6

Endpoint Behaviors (e.g., new behaviors that may be introduced in the future) is not considered an error. An unrecognized SRv6 Endpoint Behavior MUST NOT be considered invalid by the receiver, except for behaviors that involve the use of arguments (refer to Section 3.2.1 for details on argument validation). An implementation MAY log a rate-limited warning when it receives an unexpected behavior.

When multiple SRv6 SID Information Sub-TLVs are present, the ingress PE SHOULD use the SRv6 SID from the first instance of the Sub-TLV. An implementation MAY provide a local policy to override this selection.

3.2. SRv6 Service Data Sub-Sub-TLVs

The format of the SRv6 Service Data Sub-Sub-TLV is depicted below:

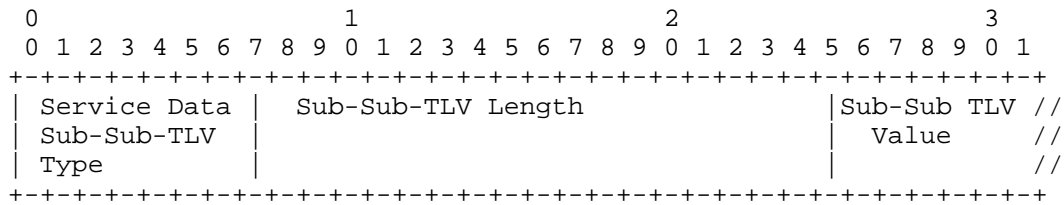


Figure 4: SRv6 Service Data Sub-Sub-TLVs

SRv6 Service Data Sub-Sub-TLV Type (1 octet):

This field identifies the type of Sub-Sub-TLV. It is assigned a value from IANA's "SRv6 Service Data Sub-Sub-TLV Types" subregistry.

SRv6 Service Data Sub-Sub-TLV Length (2 octets):

This field specifies the total length, in octets, of the Sub-Sub-TLV Value field.

SRv6 Service Data Sub-Sub-TLV Value (variable):

This field contains data specific to the Sub-Sub-TLV Type.

3.2.1. SRv6 SID Structure Sub-Sub-TLV

SRv6 Service Data Sub-Sub-TLV Type 1 is assigned for the SRv6 SID Structure Sub-Sub-TLV. The SRv6 SID Structure Sub-Sub-TLV is used to advertise the lengths of the individual parts of the SRv6 SID, as defined in [RFC8986]. The terms Locator Block and Locator Node correspond to the B and N parts, respectively, of the SRv6 Locator that is defined in Section 3.1 of [RFC8986]. It is carried as Sub-Sub-TLV in the SRv6 SID Information Sub-TLV.

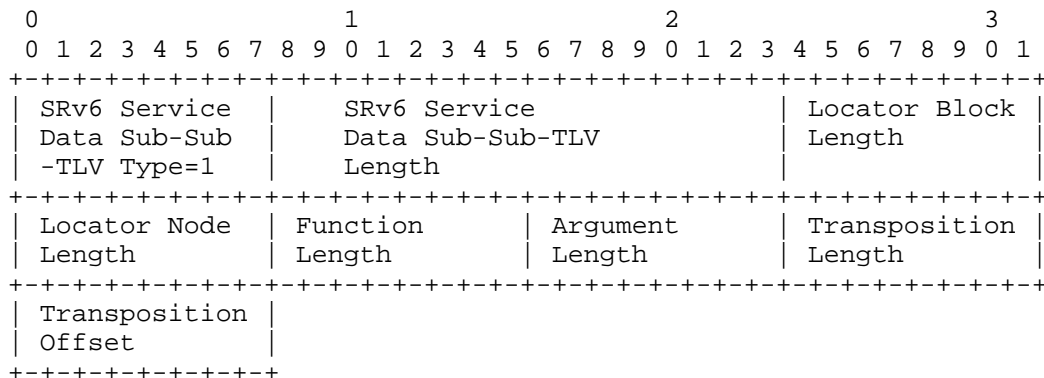


Figure 5: SRv6 SID Structure Sub-Sub-TLV

SRv6 Service Data Sub-Sub-TLV Type (1 octet):

This field is set to 1 to represent the SRv6 SID Structure Sub-

Sub-TLV.

SRv6 Service Data Sub-Sub-TLV Length (2 octets):

This field contains a total length of 6 octets.

Locator Block Length (1 octet):

This field contains the length of the SRv6 SID Locator Block in bits.

Locator Node Length (1 octet):

This field contains the length of the SRv6 SID Locator Node in bits.

Function Length (1 octet):

This field contains the length of the SRv6 SID Function in bits.

Argument Length (1 octet):

This field contains the length of the SRv6 SID Argument in bits.

Transposition Length (1 octet):

This field is the size in bits for the part of the SID that has been transposed (or shifted) into an MPLS Label field.

Transposition Offset (1 octet):

This field is the offset position in bits for the part of the SID that has been transposed (or shifted) into an MPLS Label field.

Section 4 describes mechanisms for the signaling of the SRv6 Service SID by transposing a variable part of the SRv6 SID value and carrying this variable part in existing MPLS Label fields to achieve more efficient packing of those service prefix NLRIs in BGP update messages. The SRv6 SID Structure Sub-Sub-TLV contains appropriate length fields when the SRv6 Service SID is signaled in split parts to enable the receiver to put together the SID accurately.

Transposition Offset indicates the bit position, and Transposition Length indicates the number of bits that are being taken out of the SRv6 SID value and encoded in the MPLS Label field. The bits that have been shifted out MUST be set to 0 in the SID value.

A Transposition Length of 0 indicates nothing is transposed and that the entire SRv6 SID value is encoded in the SID Information Sub-TLV. In this case, the Transposition Offset MUST be set to 0.

The size of the MPLS Label field limits the bits transposed from the SRv6 SID value into it. For example, the size of the MPLS Label field is 20 bits in [RFC4364] and [RFC8277], and the size is 24 bits in [RFC7432].

As defined in [RFC8986], the sum of the Locator Block Length (LBL), Locator Node Length (LNL), Function Length (FL), and Argument Length (AL) fields MUST be less than or equal to 128 and greater than the sum of Transposition Offset and Transposition Length.

As an example, consider that the sum of the Locator Block and the Locator Node parts is 64. For an SRv6 SID where the entire Function part of size 16 bits is transposed, the transposition offset is set to 64 and the transposition length is set to 16. While for an SRv6 SID for which the FL is 24 bits and only the lower order 20 bits are transposed (e.g., due to the limit of the MPLS Label field size), the transposition offset is set to 68 and the transposition length is set to 20.

BGP speakers that do not support this specification may misinterpret, on the reception of an SRv6-based BGP service route update, the part of the SRv6 SID encoded in an MPLS Label field(s) as MPLS label

values for MPLS-based services. Implementations supporting this specification MUST provide a mechanism to control the advertisement of SRv6-based BGP service routes on a per-neighbor and per-service basis. The details of deployment designs and implementation options are outside the scope of this document.

Arguments may be generally applicable for SIDs of only specific SRv6 Endpoint Behaviors (e.g., End.DT2M); therefore, the AL MUST be set to 0 for SIDs where the Argument is not applicable. A receiver is unable to validate the applicability of arguments for SRv6 Endpoint Behaviors that are unknown to it and hence MUST ignore SRv6 SIDs with arguments (indicated by a non-zero AL) with unknown SRv6 Endpoint Behaviors. For SIDs corresponding to an SRv6 Endpoint Behavior that is known, a receiver MUST validate that the consistency of the AL with the specific SRv6 Endpoint Behavior definition.

4. Encoding SRv6 SID Information

The SRv6 Service SID(s) for a BGP service prefix is carried in the SRv6 Services TLVs of the BGP Prefix-SID attribute.

For certain types of BGP Services, like L3VPN where a per-VRF SID allocation is used (i.e., End.DT4 or End.DT6 behaviors), the same SID is shared across multiple NLRIs, thus providing efficient packing. However, for certain other types of BGP Services, like EVPN Virtual Private Wire Service (VPWS) where a per-PW SID allocation is required (i.e., End.DX2 behavior), each NLRI would have its own unique SID, thereby resulting in inefficient packing.

To achieve efficient packing, this document allows either 1) the encoding of the SRv6 Service SID as a whole in the SRv6 Services TLVs or 2) the encoding of only the common part of the SRv6 SID (e.g., Locator) in the SRv6 Services TLVs and the encoding of the variable (e.g., Function or Argument parts) in the existing label fields specific to that service encoding. This latter form of encoding is referred to as the Transposition Scheme, where the SRv6 SID Structure Sub-Sub-TLV describes the sizes of the parts of the SRv6 SID and also indicates the offset of the variable part along with its length in the SRv6 SID value. The use of the Transposition Scheme is RECOMMENDED for the specific service encodings that allow it, as described further in Sections 5 and 6.

As an example, for the EVPN VPWS service prefix described further in Section 6.1.2, the Function part of the SRv6 SID is encoded in the MPLS Label field of the NLRI, and the SID value in the SRv6 Services TLV carries only the Locator part with the SRv6 SID Structure Sub-Sub-TLV. The SRv6 SID Structure Sub-Sub-TLV defines the lengths of Locator Block, Locator Node, and Function parts (Arguments are not applicable for the End.DX2 behavior). Transposition Offset indicates the bit position, and Transposition Length indicates the number of bits that are being taken out of the SID and put into the label field.

In yet another example, for the EVPN Ethernet Auto-Discovery (A-D) per Ethernet Segment (ES) route described further in Section 6.1.1, only the Argument of the SID needs to be signaled. This Argument part of the SRv6 SID MAY be transposed in the Ethernet Segment Identifier (ESI) Label field of the ESI Label extended community, and the SID value in the SRv6 Services TLV is set to 0 along with the inclusion of the SRv6 SID Structure Sub-Sub-TLV. The SRv6 SID Structure Sub-Sub-TLV defines the lengths of Locator Block, Locator Node, Function, and Argument parts. The offset and length of the Argument part SID value moved to the label field is set in transposition offset and length of the SID Structure TLV. The receiving router is then able to put together the entire SRv6 Service SID (e.g., for the End.DT2M behavior), placing the label value

received in the ESI Label field of the Ethernet A-D per ES route into the correct transposition offset and length in the SRv6 SID with the End.DT2M behavior received for an EVPN Route Type 3 value.

5. BGP-Based L3 Service over SRv6

BGP egress nodes (egress PEs) advertise a set of reachable prefixes. Standard BGP update propagation schemes [RFC4271], which may make use of route reflectors [RFC4456], are used to propagate these prefixes. BGP ingress nodes (ingress PEs) receive these advertisements and may add the prefix to the RIB in an appropriate VRF.

Egress PEs that support SRv6-based L3 services advertise overlay service prefixes along with a Service SID enclosed in an SRv6 L3 Service TLV within the BGP Prefix-SID attribute. This TLV serves two purposes -- first, it indicates that the egress PE supports SRv6 overlay, and the BGP ingress PE receiving this route MUST perform IPv6 encapsulation and insert an SRH [RFC8754] when required; second, it indicates the value of the Service SID to be used in the encapsulation.

Thus, the Service SID signaled only has local significance at the egress PE, where it may be allocated or configured on a per-Customer-Edge (CE) or per-VRF basis. In practice, the SID may encode a cross-connect to a specific address family table (End.DT) or next hop / interface (End.DX), as defined in [RFC8986].

The SRv6 Service SID SHOULD be routable (refer to Section 3.3 of [RFC8986]) within the Autonomous System (AS) of the egress PE and serves the dual purpose of providing reachability between ingress PE and egress PE while also encoding the SRv6 Endpoint Behavior.

When steering for SRv6 services is based on shortest path forwarding (e.g., best effort or IGP Flexible Algorithm [IGP-FLEX-ALGO]) to the egress PE, the ingress PE encapsulates the IPv4 or IPv6 customer packet in an outer IPv6 header (using H.Encaps or H.Encaps.Red flavors specified in [RFC8986]), where the destination address is the SRv6 Service SID associated with the related BGP route update. Therefore, the ingress PE MUST perform a resolvability check for the SRv6 Service SID before considering the received prefix for the BGP best path computation. The resolvability is evaluated as per [RFC4271]. If the SRv6 SID is reachable via more than one forwarding table, local policy is used to determine which table to use. The result of an SRv6 Service SID resolvability (e.g., when provided via IGP Flexible Algorithm) can be ignored if the ingress PE has a local policy that allows an alternate steering mechanism to reach the egress PE. The details of such steering mechanisms are outside the scope of this document.

For service over SRv6 core, the egress PE sets the BGP next hop to one of its IPv6 addresses. Such an address MAY be covered by the SRv6 Locator from which the SRv6 Service SID is allocated. The BGP next hop is used for tracking the reachability of the egress PE based on existing BGP procedures.

When the BGP route received at an ingress PE is colored with a Color Extended Community and a valid SRv6 Policy is available, the steering for service flows is performed as described in Section 8 of [SEGMENT-ROUTING-POLICY]. When the ingress PE determines (with the help of the SRv6 SID Structure) that the Service SID belongs to the same SRv6 Locator as the last SRv6 SID (of the egress PE) in the SR Policy segment list, it MAY exclude that last SRv6 SID when steering the service flow. For example, the effective segment list of the SRv6 Policy associated with SID list <S1, S2, S3> would be <S1, S2, S3-Service-SID>.

5.1. IPv4 VPN over SRv6 Core

The MP_REACH_NLRI over SRv6 core is encoded according to IPv4 VPN unicast over IPv6 core defined in [RFC8950].

The label field of IPv4-VPN NLRI is encoded as specified in [RFC8277] with the 20-bit Label Value set to the whole or a portion of the Function part of the SRv6 SID when the Transposition Scheme of encoding (Section 4) is used; otherwise, it is set to Implicit NULL. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 20 and less than or equal to the FL.

The SRv6 Service SID is encoded as part of the SRv6 L3 Service TLV. The SRv6 Endpoint Behavior SHOULD be one of these: End.DX4, End.DT4, or End.DT46.

5.2. IPv6 VPN over SRv6 Core

The MP_REACH_NLRI over SRv6 core is encoded according to IPv6 VPN over IPv6 core, as defined in [RFC4659].

The label field of the IPv6-VPN NLRI is encoded as specified in [RFC8277] with the 20-bit Label Value set to the whole or a portion of the Function part of the SRv6 SID when the Transposition Scheme of encoding (Section 4) is used; otherwise, it is set to Implicit NULL. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 20 and less than or equal to the FL.

The SRv6 Service SID is encoded as part of the SRv6 L3 Service TLV. The SRv6 Endpoint Behavior SHOULD be one of these: End.DX6, End.DT6, or End.DT46.

5.3. Global IPv4 over SRv6 Core

The MP_REACH_NLRI over SRv6 core is encoded according to IPv4 over IPv6 core, as defined in [RFC8950].

SRv6 Service SID is encoded as part of the SRv6 L3 Service TLV. The SRv6 Endpoint Behavior SHOULD be one of these: End.DX4, End.DT4, or End.DT46.

5.4. Global IPv6 over SRv6 Core

The MP_REACH_NLRI over SRv6 core is encoded according to [RFC2545].

The SRv6 Service SID is encoded as part of the SRv6 L3 Service TLV. The SRv6 Endpoint Behavior SHOULD be one of these: End.DX6, End.DT6, or End.DT46.

6. BGP-Based Ethernet VPN (EVPN) over SRv6

[RFC7432] provides an extendable method of building an EVPN overlay. It primarily focuses on MPLS-based EVPNs, and [RFC8365] extends to IP-based EVPN overlays. [RFC7432] defines Route Types 1, 2, and 3, which carry prefixes and MPLS Label fields; the Label fields have a specific use for MPLS encapsulation of EVPN traffic. Route Type 5 carrying MPLS label information (and thus encapsulation information) for an EVPN is defined in [RFC9136]. Route Types 6, 7, and 8 are defined in [RFC9251].

- * Ethernet Auto-Discovery (A-D) route (Route Type 1)
- * MAC/IP Advertisement route (Route Type 2)
- * Inclusive Multicast Ethernet Tag route (Route Type 3)

- * Ethernet Segment route (Route Type 4)
- * IP Prefix route (Route Type 5)
- * Selective Multicast Ethernet Tag route (Route Type 6)
- * Multicast Membership Report Synch route (Route Type 7)
- * Multicast Leave Synch route (Route Type 8)

The specifications for other EVPN Route Types are outside the scope of this document.

To support SRv6-based EVPN overlays, one or more SRv6 Service SIDs are advertised with Route Types 1, 2, 3, and 5. The SRv6 Service SID(s) per Route Type is advertised in SRv6 L3/L2 Service TLVs within the BGP Prefix-SID attribute. Signaling of the SRv6 Service SID(s) serves two purposes -- first, it indicates that the BGP egress device supports SRv6 overlay, and the BGP ingress device receiving this route MUST perform IPv6 encapsulation and insert an SRH [RFC8754] when required; second, it indicates the value of the Service SID(s) to be used in the encapsulation.

The SRv6 Service SID SHOULD be routable (refer to Section 3.3 of [RFC8986]) within the AS of the egress PE and serves the dual purpose of providing reachability between the ingress PE and egress PE while also encoding the SRv6 Endpoint Behavior.

When steering for SRv6 services is based on shortest path forwarding (e.g., best effort or IGP Flexible Algorithm [IGP-FLEX-ALGO]) to the egress PE, the ingress PE encapsulates the customer Layer 2 Ethernet packet in an outer IPv6 header (using H.Encaps.L2 or H.Encaps.L2.Red flavors specified in [RFC8986]) where the destination address is the SRv6 Service SID associated with the related BGP route update. Therefore, the ingress PE MUST perform a resolvability check for the SRv6 Service SID before considering the received prefix for the BGP best path computation. The resolvability is evaluated as per [RFC4271]. If the SRv6 SID is reachable via more than one forwarding table, local policy is used to determine which table to use. The result of an SRv6 Service SID resolvability (e.g., when provided via IGP Flexible Algorithm) can be ignored if the ingress PE has a local policy that allows an alternate steering mechanism to reach the egress PE. The details of such steering mechanisms are outside the scope of this document.

For service over SRv6 core, the egress PE sets the BGP next hop to one of its IPv6 addresses. Such an address MAY be covered by the SRv6 Locator from which the SRv6 Service SID is allocated. The BGP next hop is used for tracking the reachability of the egress PE based on existing BGP procedures.

When the BGP route received at an ingress PE is colored with a Color Extended Community and a valid SRv6 Policy is available, the steering for service flows is performed as described in Section 8 of [SEGMENT-ROUTING-POLICY]. When the ingress PE determines (with the help of the SRv6 SID Structure) that the Service SID belongs to the same SRv6 Locator as the last SRv6 SID (of the egress PE) in the SR Policy segment list, it MAY exclude that last SRv6 SID when steering the service flow. For example, the effective segment list of the SRv6 Policy associated with SID list <S1, S2, S3> would be <S1, S2, S3-Service-SID>.

6.1. Ethernet Auto-Discovery Route over SRv6 Core

Ethernet A-D routes are Route Type 1, as defined in [RFC7432], and may be used to achieve split-horizon filtering, fast convergence, and

aliasing. EVPN Route Type 1 is also used in EVPN-VPWS as well as in EVPN-flexible cross-connect, mainly to advertise point-to-point service IDs.

As a reminder, EVPN Route Type 1 is encoded as follows:

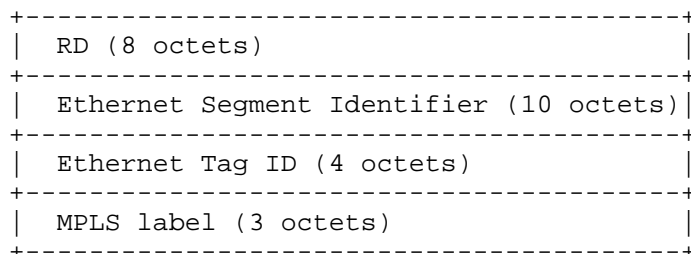


Figure 6: EVPN Route Type 1

6.1.1. Ethernet A-D per ES Route

Ethernet A-D per ES route NLRI encoding over SRv6 core is as per [RFC7432].

The 24-bit ESI Label field of the ESI Label extended community carries the whole or a portion of the Argument part of the SRv6 SID when the ESI filtering approach is used along with the Transposition Scheme of encoding (Section 4); otherwise, it is set to Implicit NULL in the higher-order 20 bits (i.e., as 0x000030). In either case, the value is set in the 24 bits. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the AL.

A Service SID enclosed in an SRv6 L2 Service TLV within the BGP Prefix-SID attribute is advertised along with the A-D route. The SRv6 Endpoint Behavior SHOULD be End.DT2M. When the ESI filtering approach is used, the Service SID is used to signal the Arg.FE2 SID Argument for applicable End.DT2M behavior [RFC8986]. When the local-bias approach [RFC8365] is used, the Service SID MAY be of value 0.

6.1.2. Ethernet A-D per EVI Route

Ethernet A-D per EVPN Instance (EVI) route NLRI encoding over SRv6 core is similar to what is described in [RFC7432] and [RFC8214] with the following change:

MPLS Label:

The 24-bit field carries the whole or a portion of the Function part of the SRv6 SID when the Transposition Scheme of encoding (Section 4) is used; otherwise, it is set to Implicit NULL in the higher-order 20 bits (i.e., as 0x000030). In either case, the value is set in the 24 bits. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the FL.

A Service SID enclosed in an SRv6 L2 Service TLV within the BGP Prefix-SID attribute is advertised along with the A-D route. The SRv6 Endpoint Behavior SHOULD be one of these: End.DX2, End.DX2V, or End.DT2U.

6.2. MAC/IP Advertisement Route over SRv6 Core

EVPN Route Type 2 is used to advertise unicast traffic Media Access Control (MAC) + IP address reachability through MP-BGP to all other PEs in a given EVPN instance.

As a reminder, EVPN Route Type 2 is encoded as follows:

RD (8 octets)
Ethernet Segment Identifier (10 octets)
Ethernet Tag ID (4 octets)
MAC Address Length (1 octet)
MAC Address (6 octets)
IP Address Length (1 octet)
IP Address (0, 4, or 16 octets)
MPLS Label1 (3 octets)
MPLS Label2 (0 or 3 octets)

Figure 7: EVPN Route Type 2

NLRI encoding over SRv6 core is similar to what is described in [RFC7432] with the following changes:

MPLS Label1:

This is associated with the SRv6 L2 Service TLV. This 24-bit field carries the whole or a portion of the Function part of the SRv6 SID when the Transposition Scheme of encoding (Section 4) is used; otherwise, it is set to Implicit NULL in the higher-order 20 bits (i.e., as 0x000030). In either case, the value is set in the 24 bits. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the FL.

MPLS Label2:

This is associated with the SRv6 L3 Service TLV. This 24-bit field carries the whole or a portion of the Function part of the SRv6 SID when the Transposition Scheme of encoding (Section 4) is used; otherwise, it is set to Implicit NULL in the higher-order 20 bits (i.e., as 0x000030). In either case, the value is set in the 24 bits. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the FL.

Service SIDs enclosed in the SRv6 L2 Service TLV and optionally in the SRv6 L3 Service TLV within the BGP Prefix-SID attribute are advertised along with the MAC/IP Advertisement route.

Described below are different types of Route Type 2 advertisements.

6.2.1. MAC/IP Advertisement Route with MAC Only

MPLS Label1:

This is associated with the SRv6 L2 Service TLV. This 24-bit field carries the whole or a portion of the Function part of the SRv6 SID when the Transposition Scheme of encoding (Section 4) is used; otherwise, it is set to Implicit NULL in the higher-order 20 bits (i.e., as 0x000030). In either case, the value is set in the 24 bits. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the FL.

A Service SID enclosed in an SRv6 L2 Service TLV within the BGP Prefix-SID attribute is advertised along with the route. The SRv6

Endpoint Behavior SHOULD be one of these: End.DX2 or End.DT2U.

6.2.2. MAC/IP Advertisement Route with MAC+IP

MPLS Label1:

This is associated with the SRv6 L2 Service TLV. This 24-bit field carries the whole or a portion of the Function part of the SRv6 SID when the Transposition Scheme of encoding (Section 4) is used; otherwise, it is set to Implicit NULL in the higher-order 20 bits (i.e., as 0x000030). In either case, the value is set in the 24 bits. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the FL.

MPLS Label2:

This is associated with the SRv6 L3 Service TLV. This 24-bit field carries the whole or a portion of the Function part of the SRv6 SID when the Transposition Scheme of encoding (Section 4) is used; otherwise, it is set to Implicit NULL in the higher-order 20 bits (i.e., as 0x000030). In either case, the value is set in the 24 bits. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the FL.

An L2 Service SID enclosed in an SRv6 L2 Service TLV within the BGP Prefix-SID attribute is advertised along with the route. In addition, an L3 Service SID enclosed in an SRv6 L3 Service TLV within the BGP Prefix-SID attribute MAY also be advertised along with the route. The SRv6 Endpoint Behavior SHOULD be one of these: for the L2 Service SID, End.DX2 or End.DT2U and for the L3 Service SID, End.DT46, End.DT4, End.DT6, End.DX4, or End.DX6.

6.3. Inclusive Multicast Ethernet Tag Route over SRv6 Core

EVPN Route Type 3 is used to advertise multicast traffic reachability information through MP-BGP to all other PEs in a given EVPN instance.

As a reminder, EVPN Route Type 3 is encoded as follows:

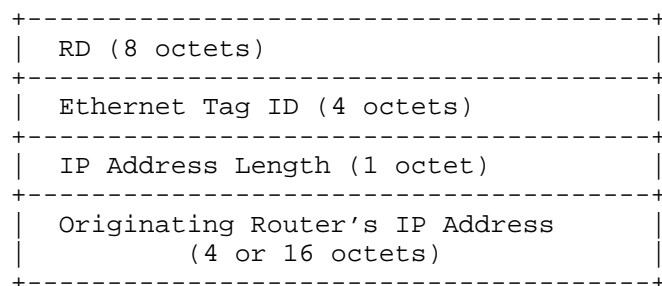
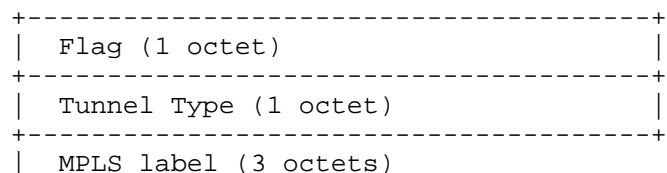


Figure 8: EVPN Route Type 3

NLRI encoding over SRv6 core is similar to what is described in [RFC7432].

The P-Multicast Service Interface (PMSI) Tunnel Attribute [RFC6514] is used to identify the Provider tunnel (P-tunnel) used for sending Broadcast, Unknown Unicast, or Multicast (BUM) traffic. The format of the PMSI Tunnel Attribute is encoded as follows over SRv6 core:



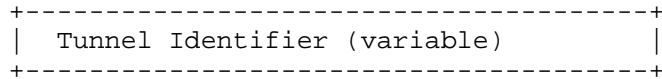


Figure 9: PMSI Tunnel Attribute

Flag:

This field has a value of 0, as defined per [RFC7432].

Tunnel Type:

This field is defined per [RFC6514].

MPLS label:

This 24-bit field carries the whole or a portion of the Function part of the SRv6 SID when ingress replication is used and the Transposition Scheme of encoding (Section 4) is used; otherwise, it is set as defined in [RFC6514]. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the FL.

Tunnel Identifier:

This field is the IP address of egress PE.

A Service SID enclosed in an SRv6 L2 Service TLV within the BGP Prefix-SID attribute is advertised along with the route. The SRv6 Endpoint Behavior SHOULD be End.DT2M.

- * When ESI-based filtering is used for multihoming or Ethernet Tree (E-Tree) procedures, the ESI Filtering Argument (the Arg.FE2 notation introduced in [RFC8986]) of the Service SID carried along with EVPN Route Type 1 SHOULD be merged with the applicable End.DT2M SID of Route Type 3 advertised by the remote PE by doing a bitwise logical-OR operation to create a single SID on the ingress PE. Details of split-horizon, ESI-based filtering mechanisms for multihoming are described in [RFC7432]. Details of filtering mechanisms for Leaf-originated BUM traffic in EVPN E-Tree services are provided in [RFC8317].
- * When "local-bias" is used as the multihoming split-horizon method, the ESI Filtering Argument SHOULD NOT be merged with the corresponding End.DT2M SID on the ingress PE. Details of the local-bias procedures are described in [RFC8365].

Usage of multicast trees as P-tunnels is outside the scope of this document.

6.4. Ethernet Segment Route over SRv6 Core

As a reminder, an Ethernet Segment route (i.e., EVPN Route Type 4) is encoded as follows:

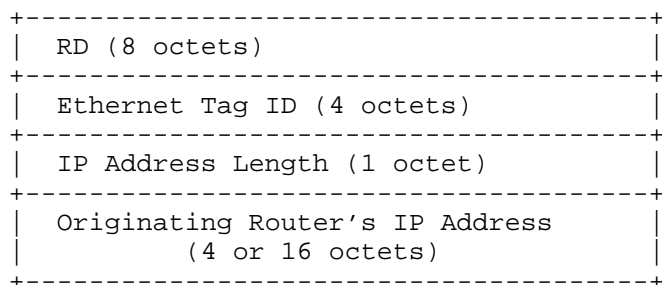


Figure 10: EVPN Route Type 4

NLRI encoding over SRv6 core is similar to what is described in [RFC7432].

SRv6 Service TLVs within the BGP Prefix-SID attribute are not advertised along with this route. The processing of the route has not changed -- it remains as described in [RFC7432].

6.5. IP Prefix Route over SRv6 Core

EVPN Route Type 5 is used to advertise IP address reachability through MP-BGP to all other PEs in a given EVPN instance. The IP address may include a host IP prefix or any specific subnet.

As a reminder, EVPN Route Type 5 is encoded as follows:

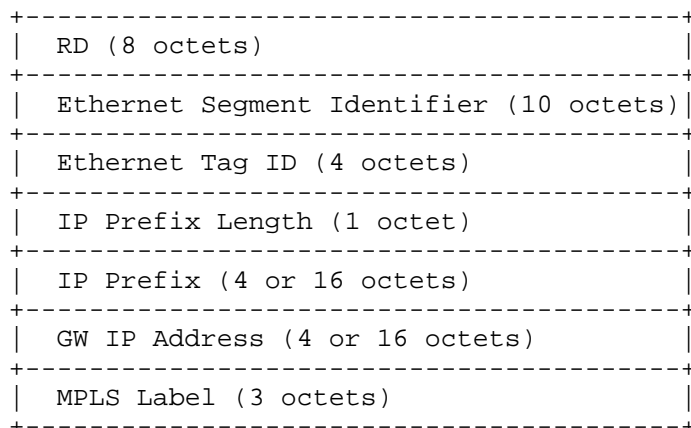


Figure 11: EVPN Route Type 5

NLRI encoding over SRv6 core is similar to what is described in [RFC9136] with the following change:

MPLS Label:

This 24-bit field carries the whole or a portion of the Function part of the SRv6 SID when the Transposition Scheme of encoding (Section 4) is used; otherwise, it is set to Implicit NULL in the higher-order 20 bits (i.e., as 0x000030). In either case, the value is set in the 24 bits. When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the FL.

The SRv6 Service SID is encoded as part of the SRv6 L3 Service TLV. The SRv6 Endpoint Behavior SHOULD be one of these: End.DT4, End.DT6, End.DT46, End.DX4, or End.DX6.

6.6. EVPN Multicast Routes (Route Types 6, 7, and 8) over SRv6 Core

These routes do not require the advertisement of SRv6 Service TLVs along with them. Similar to EVPN Route Type 4, the BGP next hop is equal to the IPv6 address of egress PE.

7. Error Handling

In case of any errors encountered while processing SRv6 Service TLVs, the details of the error SHOULD be logged for further analysis.

If multiple instances of the SRv6 L3 Service TLV are encountered, all but the first instance MUST be ignored.

If multiple instances of the SRv6 L2 Service TLV are encountered, all but the first instance MUST be ignored.

An SRv6 Service TLV is considered malformed in the following cases:

- * The TLV Length is less than 1.
- * The TLV Length is inconsistent with the length of the BGP Prefix-SID attribute.
- * At least one of the constituent Sub-TLVs is malformed.

An SRv6 Service Sub-TLV is considered malformed in the following case:

- * The Sub-TLV Length is inconsistent with the length of the enclosing SRv6 Service TLV.

An SRv6 SID Information Sub-TLV is considered malformed in the following cases:

- * The Sub-TLV Length is less than 21.
- * The Sub-TLV Length is inconsistent with the length of the enclosing SRv6 Service TLV.
- * At least one of the constituent Sub-Sub-TLVs is malformed.

An SRv6 Service Data Sub-Sub-TLV is considered malformed in the following case:

- * The Sub-Sub-TLV Length is inconsistent with the length of the enclosing SRv6 service Sub-TLV.

Any TLV, Sub-TLV, or Sub-Sub-TLV is not considered malformed because its Type is unrecognized.

Any TLV, Sub-TLV, or Sub-Sub-TLV is not considered malformed because of failing any semantic validation of its Value field.

The SRv6 overlay service requires the Service SID for forwarding. The treat-as-withdraw action [RFC7606] MUST be performed when at least one malformed SRv6 Service TLV is present in the BGP Prefix-SID attribute.

The SRv6 SID value in the SRv6 SID Information Sub-TLV is invalid when the SID Structure Sub-Sub-TLV transposition length is greater than the number of bits of the label field or if any of the conditions for the fields of the Sub-Sub-TLV, as specified in Section 3.2.1, is not met. The transposition offset and length MUST be 0 when the Sub-Sub-TLV is advertised along with routes where the Transposition Scheme is not applicable (e.g., for global IPv6 service [RFC2545] where there is no label field). The path having any such Prefix-SID attribute without any valid SRv6 SID information MUST be considered ineligible during the selection of the best path for the corresponding prefix.

8. IANA Considerations

8.1. BGP Prefix-SID TLV Types Registry

This document introduces two new TLV Types of the BGP Prefix-SID attribute. IANA has assigned Type values in the "BGP Prefix-SID TLV Types" subregistry as follows:

Value	Type	Reference
4	Deprecated	RFC 9252
5	SRv6 L3 Service TLV	RFC 9252

6	SRv6 L2 Service TLV	RFC 9252
---	---------------------	----------

Table 1: BGP Prefix-SID TLV Types
Subregistry

Value 4 previously corresponded to the SRv6-VPN SID TLV, which was specified in earlier draft versions of this document and used by early implementations of this specification. It was deprecated and replaced by the SRv6 L3 Service and SRv6 L2 Service TLVs.

8.2. SRv6 Service Sub-TLV Types Registry

IANA has created and now maintains a new subregistry called "SRv6 Service Sub-TLV Types" under the "Border Gateway Protocol (BGP) Parameters" registry. The registration procedures, per [RFC8126], for this subregistry are according to Table 2.

Range	Registration Procedures
1-127	IETF Review
128-254	First Come First Served
255	IETF Review

Table 2: SRv6 Service Sub-TLV
Types Subregistry Registration
Procedures

IANA has populated this subregistry as follows. Note that the SRv6 SID Information Sub-TLV is defined in this document:

Value	Type	Reference
0	Reserved	RFC 9252
1	SRv6 SID Information Sub-TLV	RFC 9252
255	Reserved	RFC 9252

Table 3: SRv6 Service Sub-TLV Types Subregistry
Initial Contents

8.3. SRv6 Service Data Sub-Sub-TLV Types Registry

IANA has created and now maintains a new subregistry called "SRv6 Service Data Sub-Sub-TLV Types" under the "Border Gateway Protocol (BGP) Parameters" registry. The registration procedures for this subregistry are according to Table 4.

Range	Registration Procedure
1-127	IETF Review
128-254	First Come First Served
255	IETF Review

Table 4: SRv6 Service Data Sub-
Sub-TLV Types Subregistry
Registration Procedures

The following Sub-Sub-TLV Type is defined in this document:

Value	Type	Reference
0	Reserved	RFC 9252
1	SRv6 SID Structure Sub-Sub-TLV	RFC 9252
255	Reserved	RFC 9252

Table 5: SRv6 Service Data Sub-Sub-TLV Types
Subregistry Initial Contents

8.4. BGP SRv6 Service SID Flags Registry

IANA has created and now maintains a new subregistry called "BGP SRv6 Service SID Flags" under the "Border Gateway Protocol (BGP) Parameters" registry. The registration procedure for this subregistry is IETF Review, and all 8-bit positions of the flags are currently unassigned.

8.5. SAFI Values Registry

IANA has added this document as a reference for value 128 ("MPLS-labeled VPN address") in the "SAFI Values" subregistry under the "Subsequent Address Family Identifiers (SAFI) Parameters" registry.

9. Security Considerations

This document specifies extensions to the BGP protocol for the signaling of services for SRv6. These specifications leverage existing BGP protocol mechanisms for the signaling of various types of services. It also builds upon existing elements of the SR architecture (more specifically, SRv6). As such, this section largely provides pointers (as a reminder) to the security considerations of those existing specifications while also covering certain, newer security aspects for the specifications newly introduced by this document.

9.1. Considerations Related to BGP Sessions

Techniques related to authentication of BGP sessions for securing messages between BGP peers, as discussed in the BGP specification [RFC4271] and in the security analysis for BGP [RFC4272], apply. The discussion of the use of the TCP Authentication Option to protect BGP sessions is found in [RFC5925], while [RFC6952] includes an analysis of BGP keying and authentication issues. This document does not introduce any additional BGP session security considerations.

9.2. Considerations Related to BGP Services

This document does not introduce new services or BGP NLRI types but extends the signaling of existing ones for SRv6. Therefore, the security considerations for the respective BGP services, such as BGP IPv4 over IPv6 NH [RFC8950], BGP IPv6 L3VPN [RFC4659], BGP IPv6 [RFC2545], BGP EVPN [RFC7432], and IP EVPN [RFC9136], apply as discussed in their respective documents. [RFC8669] discusses mechanisms to prevent the leaking of the BGP Prefix-SID attribute, which carries SR information, outside the SR domain.

As a reminder, several of the BGP services (i.e., the AFI/SAFI used for their signaling) were initially introduced for one encapsulation mechanism and later extended for others, e.g., EVPN MPLS [RFC7432] was extended for Virtual eXtensible Local Area Network (VXLAN) encapsulation and Network Virtualization Using Generic Routing Encapsulation (NVGRE) [RFC8365]. [RFC9012] enables the use of various IP encapsulation mechanisms along with different BGP SAFIs for their respective services. The existing filtering mechanisms for preventing the leak of the encapsulation information (carried in BGP attributes) and preventing the advertisement of prefixes from the provider's internal address space (especially the SRv6 Block, as discussed in [RFC8986]) to external peers (or into the Internet) also apply in the case of SRv6.

Specific to SRv6, a misconfiguration or error in the BGP filtering mechanisms mentioned above may result in exposing information, such as SRv6 Service SIDs to external peers or other unauthorized entities. However, an attempt to exploit this information or to raise an attack by injecting packets into the network (e.g., customer networks in case of VPN services) is mitigated by the existing SRv6 data plane security mechanisms, as described in the next section.

9.3. Considerations Related to SR over IPv6 Data Plane

This section provides a brief reminder and an overview of the security considerations related to SRv6 with pointers to existing specifications. This document introduces no new security considerations of its own from the SRv6 data plane perspective.

SRv6 operates within a trusted SR domain. The data packets corresponding to service flows between PE routers are encapsulated (using SRv6 SIDs advertised via BGP) and carried within this trusted SR domain (e.g., within a single AS or between multiple ASes within a single provider network).

The security considerations of the SR architecture are covered by [RFC8402]. More detailed security considerations, specifically of SRv6 and SRH, are covered by [RFC8754] as they relate to SR Attacks (Section 7.1), Service Theft (Section 7.2), and Topology Disclosure (Section 7.3). As such, an operator deploying SRv6 MUST follow the considerations described in Section 7 of [RFC8754] to implement the infrastructure Access Control Lists (ACLs) and the recommendations described in BCP 38 [RFC2827] and BCP 84 [RFC3704].

The SRv6 deployment and SID allocation guidelines, as described in [RFC8986], simplify the deployment of the ACL filters (e.g., a single ACL corresponding to the SRv6 Block applied to the external interfaces on border nodes is sufficient to block packets destined to any SRv6 SID in the domain from external/unauthorized networks). While there is an assumed trust model within an SR domain, such that any node sending a packet to an SRv6 SID is assumed to be allowed to do so, there is also the option of using an SRH Hashed Message Authentication Code (HMAC) TLV [RFC8754], as described in [RFC8986], for validation.

The SRv6 Endpoint Behaviors implementing the services signaled in this document are defined in [RFC8986]; hence, the security considerations of that document apply. These considerations are independent of the protocol used for service deployment, i.e., independent of BGP signaling of SRv6 services.

These considerations help protect transit traffic as well as services, such as VPNs, to avoid service theft or injection of traffic into customer VPNs.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC2545] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing", RFC 2545, DOI 10.17487/RFC2545, March 1999, <<https://www.rfc-editor.org/info/rfc2545>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, DOI 10.17487/RFC4364, February 2006, <<https://www.rfc-editor.org/info/rfc4364>>.
- [RFC4456] Bates, T., Chen, E., and R. Chandra, "BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006, <<https://www.rfc-editor.org/info/rfc4456>>.
- [RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP/MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, DOI 10.17487/RFC4659, September 2006, <<https://www.rfc-editor.org/info/rfc4659>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.
- [RFC7432] Sajassi, A., Ed., Aggarwal, R., Bitar, N., Isaac, A., Uttaro, J., Drake, J., and W. Henderickx, "BGP MPLS-Based Ethernet VPN", RFC 7432, DOI 10.17487/RFC7432, February 2015, <<https://www.rfc-editor.org/info/rfc7432>>.
- [RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.
- [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>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC8214] Boutros, S., Sajassi, A., Salam, S., Drake, J., and J. Rabadan, "Virtual Private Wire Service Support in Ethernet VPN", RFC 8214, DOI 10.17487/RFC8214, August 2017, <<https://www.rfc-editor.org/info/rfc8214>>.

- [RFC8277] Rosen, E., "Using BGP to Bind MPLS Labels to Address Prefixes", RFC 8277, DOI 10.17487/RFC8277, October 2017, <<https://www.rfc-editor.org/info/rfc8277>>.
- [RFC8317] Sajassi, A., Ed., Salam, S., Drake, J., Uttaro, J., Boutros, S., and J. Rabadan, "Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) and Provider Backbone Bridging EVPN (PBB-EVPN)", RFC 8317, DOI 10.17487/RFC8317, January 2018, <<https://www.rfc-editor.org/info/rfc8317>>.
- [RFC8365] Sajassi, A., Ed., Drake, J., Ed., Bitar, N., Shekhar, R., Uttaro, J., and W. Henderickx, "A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)", RFC 8365, DOI 10.17487/RFC8365, March 2018, <<https://www.rfc-editor.org/info/rfc8365>>.
- [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>>.
- [RFC8669] Previdi, S., Filsfils, C., Lindem, A., Ed., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", RFC 8669, DOI 10.17487/RFC8669, December 2019, <<https://www.rfc-editor.org/info/rfc8669>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8950] Litkowski, S., Agrawal, S., Ananthamurthy, K., and K. Patel, "Advertising IPv4 Network Layer Reachability Information (NLRI) with an IPv6 Next Hop", RFC 8950, DOI 10.17487/RFC8950, November 2020, <<https://www.rfc-editor.org/info/rfc8950>>.
- [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>>.
- [RFC9136] Rabadan, J., Ed., Henderickx, W., Drake, J., Lin, W., and A. Sajassi, "IP Prefix Advertisement in Ethernet VPN (EVPN)", RFC 9136, DOI 10.17487/RFC9136, October 2021, <<https://www.rfc-editor.org/info/rfc9136>>.
- [RFC9251] Sajassi, A., Thoria, S., Mishra, M., Patel, K., Drake, J., and W. Lin, "Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Proxies for Ethernet VPN (EVPN)", RFC 9251, DOI 10.17487/RFC9251, June 2022, <<https://www.rfc-editor.org/info/rfc9251>>.

10.2. Informative References

- [IGP-FLEX-ALGO] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", Work in Progress, Internet-Draft, draft-ietf-lsr-flex-algo-20, 18 May 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-flex-algo-20>>.
- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source

Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <<https://www.rfc-editor.org/info/rfc2827>>.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March 2004, <<https://www.rfc-editor.org/info/rfc3704>>.

[RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <<https://www.rfc-editor.org/info/rfc4272>>.

[RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, DOI 10.17487/RFC5925, June 2010, <<https://www.rfc-editor.org/info/rfc5925>>.

[RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.

[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>>.

[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>>.

[RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.

[SEGMENT-ROUTING-POLICY]

Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, draft-ietf-spring-segment-routing-policy-22, 22 March 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-segment-routing-policy-22>>.

Acknowledgements

The authors of this document would like to thank Stephane Litkowski, Rishabh Parekh, Xiejingrong, Rajesh M., Mustapha Aissaoui, Alexander Vainshtein, Eduard Metz, Shraddha Hegde, Eduard Vasilenko, Ron Bonica, and Joel Halpern for their comments and review of this document. The authors would also like to thank Document Shepherd Matthew Bocci for his review and AD Martin Vigoureux for his review that resulted in helpful comments for improving this document.

Contributors

Clarence Filsfils
Cisco
Email: cfilsfil@cisco.com

Satoru Matsushima
SoftBank
Email: satoru.matsushima@g.softbank.co.jp

Dirk Steinberg

Steinberg Consulting
Email: dirk@lapishills.com

Daniel Bernier
Bell Canada
Email: daniel.bernier@bell.ca

Daniel Voyer
Bell Canada
Email: daniel.voyer@bell.ca

Jonn Leddy
Individual
Email: john@leddy.net

Swadesh Agrawal
Cisco
Email: swaagraw@cisco.com

Patrice Brissette
Cisco
Email: pbrisset@cisco.com

Ali Sajassi
Cisco
Email: sajassi@cisco.com

Bart Peirens
Proximus
Belgium
Email: bart.peirens@proximus.com

Darren Dukes
Cisco
Email: ddukes@cisco.com

Pablo Camarilo
Cisco
Email: pcamaril@cisco.com

Shyam Sethuram
Cisco
Email: shyam.ioml@gmail.com

Zafar Ali
Cisco
Email: zali@cisco.com

Authors' Addresses

Gaurav Dawra (editor)
LinkedIn
United States of America
Email: gdawra.ietf@gmail.com

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

Robert Raszuk
NTT Network Innovations
940 Stewart Dr.
Sunnyvale, CA 94085
United States of America
Email: robert@raszuk.net

Bruno Decraene
Orange
France
Email: bruno.decraene@orange.com

Shunwan Zhuang
Huawei Technologies
China
Email: zhuangshunwan@huawei.com

Jorge Rabadan
Nokia
United States of America
Email: jorge.rabadan@nokia.com