

BESS Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 20, 2022

G. Dawra, Ed.
LinkedIn
C. Filsfils
K. Talaulikar, Ed.
Cisco Systems
R. Raszuk
NTT Network Innovations
B. Decraene
Orange
S. Zhuang
Huawei Technologies
J. Rabadan
Nokia
March 19, 2022

SRv6 BGP based Overlay Services
draft-ietf-bess-srv6-services-13

Abstract

This document defines procedures and messages for SRv6-based BGP services including L3VPN, EVPN, and Internet services. It builds on [RFC4364](#) "BGP/MPLS IP Virtual Private Networks (VPNs)" and [RFC7432](#) "BGP MPLS-Based Ethernet VPN".

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 September 20, 2022.

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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	4
2.	SRv6 Services TLVs	4
3.	SRv6 Service Sub-TLVs	5
3.1.	SRv6 SID Information Sub-TLV	6
3.2.	SRv6 Service Data Sub-Sub-TLVs	7
3.2.1.	SRv6 SID Structure Sub-Sub-TLV	8
4.	Encoding SRv6 SID Information	10
5.	BGP based L3 Service over SRv6	11
5.1.	IPv4 VPN Over SRv6 Core	13
5.2.	IPv6 VPN Over SRv6 Core	13
5.3.	Global IPv4 over SRv6 Core	13
5.4.	Global IPv6 over SRv6 Core	13
6.	BGP based Ethernet VPN (EVPN) over SRv6	14
6.1.	Ethernet Auto-discovery Route over SRv6 Core	15
6.1.1.	Ethernet A-D per ES Route	16
6.1.2.	Ethernet A-D per EVI Route	16
6.2.	MAC/IP Advertisement Route over SRv6 Core	17
6.2.1.	MAC/IP Advertisement Route with MAC Only	18
6.2.2.	MAC/IP Advertisement Route with MAC+IP	18
6.3.	Inclusive Multicast Ethernet Tag Route over SRv6 Core	19
6.4.	Ethernet Segment Route over SRv6 Core	21
6.5.	IP Prefix Route over SRv6 Core	21
6.6.	EVPN Multicast Routes (Route Types 6, 7, 8) over SRv6 Core	22
7.	Implementation Status	22
8.	Error Handling	23
9.	IANA Considerations	24
9.1.	BGP Prefix-SID TLV Types Registry	24
9.2.	SRv6 Service Sub-TLV Types Registry	24
9.3.	SRv6 Service Data Sub-Sub-TLV Types Registry	25
9.4.	BGP SRv6 Service SID Flags Registry	25
10.	Security Considerations	26
10.1.	BGP Session Related Considerations	26
10.2.	BGP Services Related Considerations	26
10.3.	SR over IPv6 Data Plane Related Considerations	27

11.	Acknowledgments	28
12.	Contributors	28
13.	References	29
13.1.	Normative References	29
13.2.	Informative References	32
	Authors' Addresses	33

[1.](#) Introduction

SRv6 refers to Segment Routing instantiated on the IPv6 dataplane [[RFC8402](#)].

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

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

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 Provider Edge (PE) router, such as (but not limited to), End.DT (Table lookup in a VRF) or End.DX (cross-connect to a nexthop) behaviors in the case of Layer-3 Virtual Private Network (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 Provider Edge (PE). The underlay between the PEs only needs to support plain IPv6 forwarding [[RFC8200](#)].

To provide SRv6 service in conjunction with an underlay SLA from the ingress PE to the egress PE, the egress PE colors the overlay service route with a Color Extended Community [[I-D.ietf-idr-segment-routing-te-policy](#)] for steering of flows for those routes as specified in section 8 of [[I-D.ietf-spring-segment-routing-policy](#)]. The ingress PE encapsulates the payload packet in an outer IPv6 header with the segment list of SR policy 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 SID's are part of the SRH segment list MUST support 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.

- o 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](#)] [[RFC9136](#)]. Some SRv6 Endpoint behaviors which may be encoded, but not limited to, are End.DX4, End.DT4, End.DX6, End.DT6, and End.DT46.
- o 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 Label1 for Ethernet VPN (EVPN) Route-Types as defined in [[RFC7432](#)]. Some SRv6 Endpoint behaviors which may be encoded, but not limited to, are End.DX2, End.DX2V, End.DT2U, and End.DT2M.

When an egress PE is enabled for BGP Services over SRv6 data-plane, it signals one or more SRv6 Service SIDs enclosed in SRv6 Service TLV(s) within the BGP Prefix-SID Attribute attached to MP-BGP NLRIs defined in [[RFC4760](#)] [[RFC4659](#)] [[RFC8950](#)] [[RFC7432](#)] [[RFC4364](#)] [[RFC9136](#)] where applicable as described in [Section 5](#) and [Section 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:




Figure 1: SRv6 Service TLVs

- o TLV Type (1 octet): This field is assigned values from the IANA registry "BGP Prefix-SID TLV Types". It is set to 5 for SRv6 L3 Service TLV. It is set to 6 for SRv6 L2 Service TLV.
- o TLV Length (2 octets): Specifies the total length, in octets, of the TLV Value.
- o RESERVED (1 octet): This field is reserved; it MUST be set to 0 by the sender and ignored by the receiver.
- o 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 BGP Prefix-SID Attribute with one or more SRv6 Service TLVs observes the following rules when advertising the received route to other peers:

- o if the nexthop 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 or Sub-TLV or Sub-Sub-TLV MUST be propagated unchanged.
- o if the nexthop is changed, the TLVs, Sub-TLVs, and Sub-Sub-TLVs SHOULD be updated with the locally allocated SRv6 SID information. Any unrecognized received Sub-TLVs and Sub-Sub-TLVs MUST be removed.

3. SRv6 Service Sub-TLVs

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

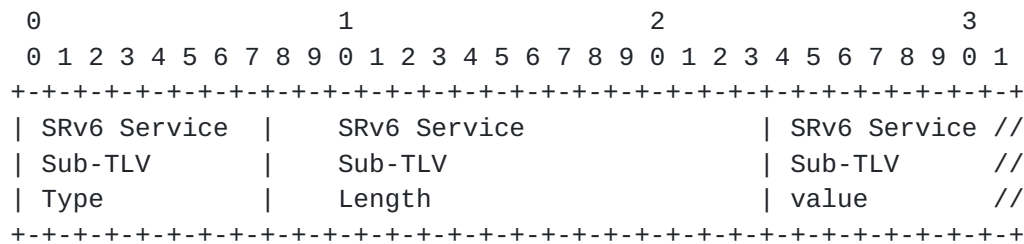


Figure 2: SRv6 Service Sub-TLVs

- o SRv6 Service Sub-TLV Type (1 octet): Identifies the type of SRv6 service information. It is assigned values from the IANA Registry "SRv6 Service Sub-TLV Types".
- o SRv6 Service Sub-TLV Length (2 octets): Specifies the total length, in octets, of the Sub-TLV Value field.
- o SRv6 Service Sub-TLV Value (variable): 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 SRv6 SID Information Sub-TLV. This Sub-TLV contains a single SRv6 SID along with its properties. Its encoding is depicted below:

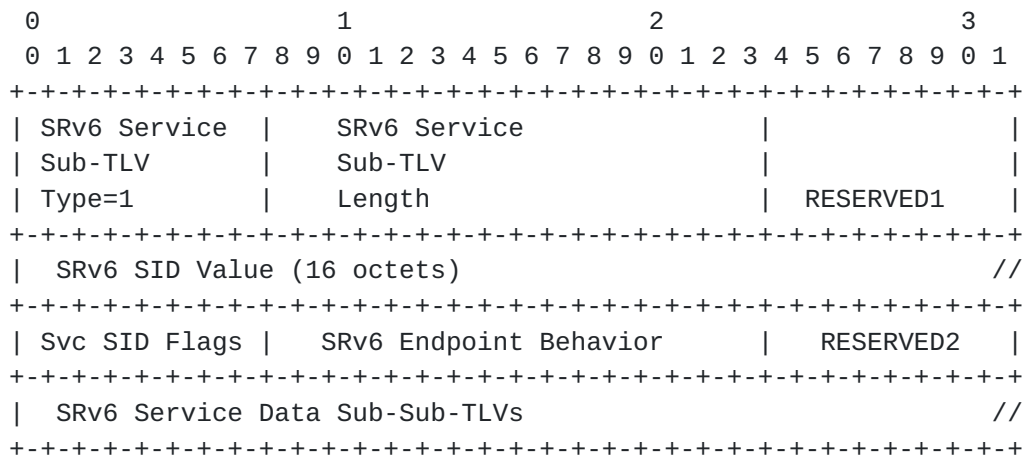


Figure 3: SRv6 SID Information Sub-TLV

- o SRv6 Service Sub-TLV Type (1 octet): This field is set to 1 to represent SRv6 SID Information Sub-TLV.

- o SRv6 Service Sub-TLV Length (2 octets): This field contains the total length, in octets, of the Value field of the Sub-TLV.
- o RESERVED1 (1 octet): MUST be set to 0 by the sender and ignored by the receiver.
- o SRv6 SID Value (16 octets): Encodes an SRv6 SID as defined in [\[RFC8986\]](#)
- o SRv6 Service SID Flags (1 octet): Encodes SRv6 Service SID Flags - none are currently defined. SHOULD be set to 0 by the sender and any unknown flags MUST be ignored by the receiver.
- o SRv6 Endpoint Behavior (2 octets): Encodes SRv6 Endpoint behavior codepoint value that is associated with SRv6 SID. The codepoints used are from the "SRv6 Endpoint Behavior" registry under the IANA "Segment Routing" parameters registry that was introduced by [\[RFC8986\]](#). The opaque endpoint behavior (i.e., value 0xFFFF) MAY be used when the advertising router wishes to abstract the actual behavior of it's locally instantiated SRv6 SID.
- o RESERVED2 (1 octet): MUST be set to 0 by the sender and ignored by the receiver.
- o SRv6 Service Data Sub-Sub-TLV Value (variable): 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 [Section 5](#) and [Section 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 endpoint behavior MUST NOT be considered invalid by the receiver. 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:

- o SRv6 Service Data Sub-Sub-TLV Length (2 octets): This field contains a total length of 6 octets.
- o Locator Block Length (1 octet): Contains the length of SRv6 SID Locator Block in bits.
- o Locator Node Length (1 octet): Contains the length of SRv6 SID Locator Node in bits.
- o Function Length (1 octet): Contains the length of SRv6 SID Function in bits.
- o Argument Length (1 octet): Contains the length of SRv6 SID Argument in bits.
- o Transposition Length (1 octet): Size in bits for the part of SID that has been transposed (or shifted) into a MPLS label field
- o Transposition Offset (1 octet): The offset position in bits for the part of SID that has been transposed (or shifted) into a MPLS label field.

[Section 4](#) describes mechanisms for signaling of the SRv6 Service SID by transposing a variable part of the SRv6 SID value and carrying them 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 put into high order bits of MPLS label field. The bits that have been shifted out MUST be set to 0 in the SID value.

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. E.g., the size of MPLS label field in [\[RFC4364\]](#) [\[RFC8277\]](#) is 20 bits while in [\[RFC7432\]](#) is 24 bits.

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, then the transposition offset is set to 64 and the transposition length is set to 16. While for an SRv6 SID where the Function length is 24 bits and only the lower order 20 bits are transposed (e.g. due to the limit of the MPLS label field size), then 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 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) and therefore the Argument length 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 non-zero argument length) with unknown endpoint behaviors. For SIDs corresponding to an endpoint behavior that is known, a receiver MUST validate that the consistency of the argument length with the specific endpoint behavior definition.

4. Encoding SRv6 SID Information

The SRv6 Service SID(s) for a BGP Service Prefix are 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 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 the encoding of the SRv6 Service SID either as a whole in the SRv6 Services TLVs or the encoding of only the common part of the SRv6 SID (e.g., Locator) in the SRv6 Services TLVs and encoding the variable (e.g., Function or Argument parts) in the existing label fields specific to that service encoding. This later 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 SRv6 SID value. The use of the Transposition Scheme is RECOMMENDED for the specific service encodings that allow it as described further in [Section 5](#) and [Section 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 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 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 label field is set in transposition offset and length of 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 which supports SRv6 based L3 services advertises 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.

The Service SID thus signaled only has local significance at the egress PE, where it may be allocated or configured on a per-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 [section 3.3 of \[RFC8986\]](#)) within the 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 [[I-D.ietf-lsr-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 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 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 next-hop is used for tracking the reachability of the egress PE based on existing BGP procedures.

When the BGP route is 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 [[I-D.ietf-spring-segment-routing-policy](#)]. When the ingress PE determines (with the help of 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 Over IPv6 Core defined in [[RFC8950](#)].

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 and otherwise 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 Function Length.

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, 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 is defined in [[RFC4659](#)].

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 and otherwise 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 Function Length.

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, End.DT46.

5.3. Global IPv4 over SRv6 Core

The MP_REACH_NLRI over SRv6 core is encoded according to IPv4 over IPv6 Core is 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, End.DT46.

5.4. Global IPv6 over SRv6 Core

The MP_REACH_NLRI over SRv6 core is encoded according to [[RFC2545](#)]

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, End.DT46.

6. BGP based Ethernet VPN (EVPN) over SRv6

[RFC7432] provides an extendable method of building an Ethernet VPN (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 EVPN is defined in [RFC9136]. Route Types 6, 7, and 8 are defined in [I-D.ietf-bess-evpn-igmp-mld-proxy].

- o Ethernet Auto-discovery Route (Route Type 1)
- o MAC/IP Advertisement Route (Route Type 2)
- o Inclusive Multicast Ethernet Tag Route (Route Type 3)
- o Ethernet Segment route (Route Type 4)
- o IP prefix route (Route Type 5)
- o Selective Multicast Ethernet Tag route (Route Type 6)
- o Multicast Membership Report Synch route (Route Type 7)
- o 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 Type 1, 2, 3, and 5. The SRv6 Service SID(s) per Route Type are advertised in SRv6 L3/L2 Service TLVs within the BGP Prefix-SID Attribute. Signaling of 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 [section 3.3 of \[RFC8986\]](#)) within the 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 [[I-D.ietf-lsr-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 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 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 next-hop is used for tracking the reachability of the egress PE based on existing BGP procedures.

When the BGP route is 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 [[I-D.ietf-spring-segment-routing-policy](#)]. When the ingress PE determines (with the help of 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 Auto-Discovery (A-D) routes are Route Type 1 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 used to advertise point-to-point services ID.

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

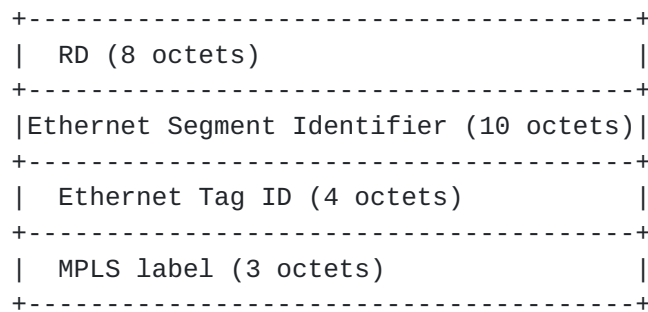


Figure 6: EVPN Route Type 1

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](#)) and otherwise set to Implicit NULL value. In either case, the value is set in the high order 20 bits (e.g., as 0x000030 in the case of Implicit NULL). When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the Argument Length.

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 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 EVI route NLRI encoding over SRv6 core is similar to [\[RFC7432\]](#) and [\[RFC8214\]](#) with the following change:

- o MPLS Label: 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 and otherwise set to Implicit NULL value. In either case, the value is set in the high order 20 bits (e.g., as 0x000030 in the case of Implicit NULL). When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the Function Length.

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, End.DT2U.

6.2. MAC/IP Advertisement Route over SRv6 Core

EVPN Route Type 2 is used to advertise unicast traffic 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 [\[RFC7432\]](#) with the following changes:

- o MPLS Label1: 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 and otherwise set to Implicit NULL value. In either case, the value is set in the high order 20 bits (e.g., as 0x000030 in the case of Implicit NULL). When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the Function Length.

- o MPLS Label2: 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 and otherwise set to Implicit NULL value. In either case, the value is set in the high order 20 bits (e.g., as 0x0000030 in the case of Implicit NULL). When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the Function Length.

Service SIDs enclosed in SRv6 L2 Service TLV and optionally in SRv6 L3 Service TLV within the BGP Prefix-SID attribute is 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

- o MPLS Label1: 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 and otherwise set to Implicit NULL value. In either case, the value is set in the high order 20 bits (e.g., as 0x0000030 in the case of Implicit NULL). When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the Function Length.

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, End.DT2U.

6.2.2. MAC/IP Advertisement Route with MAC+IP

- o MPLS Label1: 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 and otherwise set to Implicit NULL value. In either case, the value is set in the high order 20 bits (e.g., as 0x0000030 in the case of Implicit NULL). When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the Function Length.
- o MPLS Label2: 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 and otherwise set to Implicit NULL value. In either case, the value is set in the high order 20 bits (e.g., as 0x0000030 in the case of Implicit NULL). When using the

Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the Function Length.

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, End.DT2U; for the L3 Service SID - End.DT46, End.DT4, End.DT6, End.DX4, 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:

```

+-----+
|  RD (8 octets)                |
+-----+
|  Ethernet Tag ID (4 octets)   |
+-----+
|  IP Address Length (1 octet)  |
+-----+
|  Originating Router's IP Address |
|          (4 or 16 octets)      |
+-----+

```

Figure 8: EVPN Route Type 3

NLRI encoding over SRv6 core is similar to [\[RFC7432\]](#).

PMSI Tunnel Attribute [\[RFC6514\]](#) is used to identify the P-tunnel used for sending broadcast, unknown unicast, or multicast (BUM) traffic. The format of PMSI Tunnel Attribute is encoded as follows over SRv6 Core:

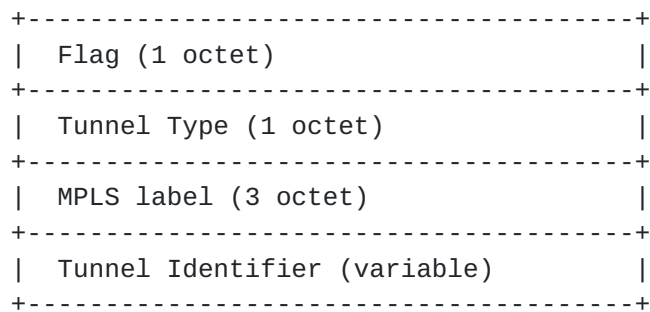


Figure 9: PMSI Tunnel Attribute

- o Flag: zero value defined per [\[RFC7432\]](#)
- o Tunnel Type: defined per [\[RFC6514\]](#)
- o 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 and 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 Function Length.
- o Tunnel Identifier: 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.

- o When ESI-based filtering is used for Multi-Homing or 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 route SHOULD be merged with the applicable End.DT2M SID of Type 3 route advertised by remote PE by doing a bit-wise 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 [\[RFC8317\]](#).
- o When "local-bias" is used as the Multi-Homing 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:

```
+-----+
|  RD (8 octets)  |
+-----+
| Ethernet Tag ID (4 octets) |
+-----+
| IP Address Length (1 octet) |
+-----+
| Originating Router's IP Address |
|          (4 or 16 octets)       |
+-----+
```

Figure 10: EVPN Route Type 4

NLRI encoding over SRv6 core is similar to [\[RFC7432\]](#).

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:

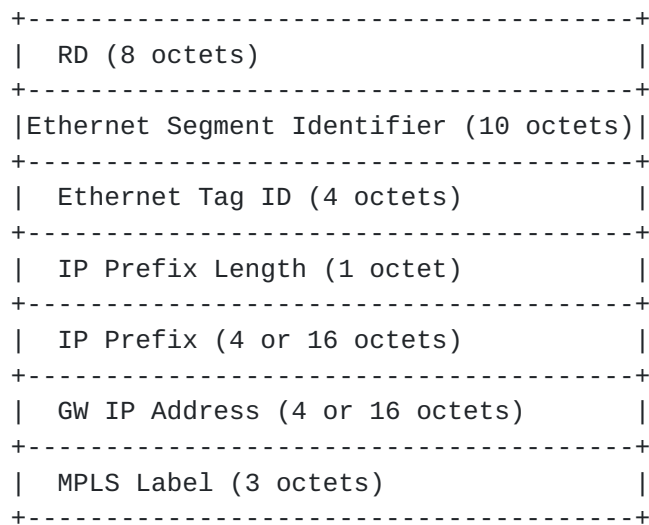


Figure 11: EVPN Route Type 5

NLRI encoding over SRv6 core is similar to [\[RFC9136\]](#) with the following change:

- o 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 and otherwise set to Implicit NULL value. In either case, the value is set in the high order 20 bits (e.g., as 0x000030 in the case of Implicit NULL). When using the Transposition Scheme, the Transposition Length MUST be less than or equal to 24 and less than or equal to the Function Length.

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, End.DX6.

6.6. EVPN Multicast Routes (Route Types 6, 7, 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 Nexthop is equal to the IPv6 address of egress PE.

7. Implementation Status

[Note to RFC Editor: This section needs to be removed before publication as RFC.]

The [\[I-D.matsushima-spring-srv6-deployment-status\]](#) describes the current deployment and implementation status of SRv6 which also includes the BGP services over SRv6 as specified in this document.

8. 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 SRv6 L3 Service TLV are encountered, all but the first instance MUST be ignored.

If multiple instances of 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:

- o the TLV Length is less than 1
- o the TLV Length is inconsistent with the length of BGP Prefix-SID attribute
- o at least one of the constituent Sub-TLVs is malformed

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

- o 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 cases:

- o the Sub-Sub-TLV Length is inconsistent with the length of the enclosing SRv6 service Sub-TLV

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

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

SRv6 overlay service requires 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.

SRv6 SID value in SRv6 SID Information Sub-TLV is invalid when 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 transposition scheme is not applicable (e.g., for Global IPv6 Service [[RFC2545](#)] where there is no label field). The path having 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.

9. IANA Considerations

9.1. 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 registry "BGP Prefix-SID TLV Types" as follows:

Value	Type	Reference
-----	-----	-----
4	Deprecated	<this document>
5	SRv6 L3 Service TLV	<this document>
6	SRv6 L2 Service TLV	<this document>

Figure 12: BGP Prefix-SID TLV Types

The value 4 previously corresponded to the SRv6-VPN SID TLV, which was specified in previous 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.

9.2. SRv6 Service Sub-TLV Types Registry

IANA is requested to create and maintain a new registry called "SRv6 Service Sub-TLV Types" under the "Border Gateway Protocol (BGP) Parameters" registry. The allocation policy for this registry is:

0 : Reserved
 1-127 : IETF Review
 128-254 : First Come First Served
 255 : Reserved

Figure 13: SRv6 Service Sub-TLV Types Allocation Policy

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

Value	Type	Reference

1	SRv6 SID Information Sub-TLV	<this document>

Figure 14: SRv6 Service Sub-TLV Types

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

IANA is requested to create and maintain a new registry called "SRv6 Service Data Sub-Sub-TLV Types" under the "Border Gateway Protocol (BGP) Parameters" registry. The allocation policy for this registry is:

0 : Reserved
 1-127 : IETF Review
 128-254 : First Come First Served
 255 : Reserved

Figure 15: SRv6 Service Data Sub-Sub-TLV Types Allocation Policy

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

Value	Type	Reference

1	SRv6 SID Structure Sub-Sub-TLV	<this document>

Figure 16: SRv6 Service Data Sub-Sub-TLV Types

9.4. BGP SRv6 Service SID Flags Registry

IANA is requested to create and maintain a new registry called "BGP SRv6 Service SID Flags" under the "Border Gateway Protocol (BGP) Parameters" registry. The allocation policy for this registry is IETF Review and all 8 bit positions of the flags are currently unassigned.

10. Security Considerations

This document specifies extensions to the BGP protocol for 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.

10.1. BGP Session Related Considerations

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.

10.2. BGP Services Related Considerations

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 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 leaking of BGP Prefix-SID attribute, that 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 VXLAN/NVGRE encapsulation [[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 to prevent 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 misconfig or error in the above mentioned BGP filtering mechanisms 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.

10.3. SR over IPv6 Data Plane Related Considerations

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 Segment Routing 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 [\[RFC8754\] section 7](#) to implement the infrastructure ACLs, [BCP 38](#) [\[RFC2827\]](#) and [BCP 84](#) [\[RFC3704\]](#) recommendations.

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 a SR domain such that any node sending packet to an SRv6 SID is assumed to be allowed to do so, there is also the option of using SRH HMAC TLV [\[RFC8754\]](#) as described in [\[RFC8986\]](#) for validation.

The SRv6 SID Endpoint behaviors implementing the services signalled in this document are defined in [\[RFC8986\]](#) and 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 VPN.

11. Acknowledgments

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, and Ron Bonica for their comments and review of this document. The authors would also like to thank Matthew Bocci for his document shepherd review and Martin Vigoureux for his AD review that resulted in helpful comments for improving this document.

12. Contributors

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

13. References

13.1. Normative References

- [I-D.ietf-bess-evpn-igmp-ml-d-proxy]
Sajassi, A., Thoria, S., Mishra, M., Drake, J., and W.
Lin, "IGMP and MLD Proxy for EVPN", [draft-ietf-bess-evpn-igmp-ml-d-proxy-19](#) (work in progress), March 2022.
- [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>>.

13.2. 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", [draft-ietf-idr-segment-routing-te-policy-16](#) (work in progress), March 2022.
- [I-D.ietf-lsr-flex-algo]
Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", [draft-ietf-lsr-flex-algo-18](#) (work in progress), October 2021.
- [I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-21](#) (work in progress), March 2022.
- [I-D.matsushima-spring-srv6-deployment-status]
Matsushima, S., Filsfils, C., Ali, Z., Li, Z., Rajaraman, K., and A. Dhamija, "SRv6 Implementation and Deployment Status", [draft-matsushima-spring-srv6-deployment-status-12](#) (work in progress), February 2022.
- [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>>.
- [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>>.

Authors' Addresses

Gaurav Dawra (editor)
LinkedIn
USA

Email: gdawra.ietf@gmail.com

Clarence Filsfils
Cisco Systems
Belgium

Email: cfilsfil@cisco.com

Ketan Talaulikar (editor)
Cisco Systems
India

Email: ketant.ietf@gmail.com

Robert Raszuk
NTT Network Innovations
940 Stewart Dr
Sunnyvale, CA 94085
USA

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
USA

Email: jorge.rabadan@nokia.com

