Inter-Domain Routing
Internet-Draft
Intended status: Standards Track
Expires: June 29, 2018

G. Dawra, Ed.
C. Filsfils
D. Dukes
P. Brissette
P. Camarilo
Cisco Systems
J. Leddy
Comcast
D. Voyer
D. Bernier
Bell Canada
D. Steinberg
Steinberg Consulting
R. Raszuk
Bloomberg LP
B. Decraene
Orange
S. Matsushima
SoftBank
December 26, 2017

# BGP Signaling of IPv6-Segment-Routing-based VPN Networks
**draft-dawra-idr-srv6-vpn-03**

## Abstract

This draft defines procedures and messages for BGP SRv6-based L3VPN and EVPN. It builds on RFC4364 "BGP/MPLS IP Virtual Private Networks (VPNs)" and RFC7432 "BGP MPLS-Based Ethernet VPN" and provides a migration path from MPLS-based VPNs to SRv6 based VPNs.

## Status of This Memo

This Internet-Draft will expire on June 29, 2018.

Copyright Notice



Table of Contents

## 1. Introduction

SRv6 refers to Segment Routing instantiated on the IPv6 dataplane [I-D.filsfils-spring-srv6-network-programming][I-D.ietf-6man-segment-routing-header].

SRv6-based VPN (SRv6-VPN) refers to the creation of VPN between PE's leveraging the SRv6 dataplane and more specifically the END.DT* (crossconnect to a VRF) and END.DX* (crossconnect to a nexthop). SRv6-L3VPN refers to the creation of Layer3 VPN service between PE's supporting an SRv6 data plane. SRv6-EVPN refers to the creation of Layer2/Layer3 VPN service between PE's supporting an SRv6 data plane.

SRv6 SID refers to a SRv6 Segment Identifier as defined in [I-D.filsfils-spring-srv6-network-programming].

SRv6-VPN SID refers to an SRv6 SID that MAY be associated with one of the END.DT or END.DX functions as defined in [I-D.filsfils-spring-srv6-network-programming].

To provide SRv6-VPN service with best-effort connectivity, the egress PE signals an SRv6-VPN SID with the VPN route. The ingress PE encapsulates the VPN packet in an outer IPv6 header where the destination address is the SRv6-VPN SID provided by the egress PE. The underlay between the PE's only need to support plain IPv6 forwarding [RFC2460].

To provide SRv6-VPN service in conjunction with an underlay SLA from the ingress PE to the egress PE, the egress PE colors the overlay VPN route with a color extended community. The ingress PE encapsulates the VPN packet in an outer IPv6 header with an SRH that contains the SR policy associated with the related SLA followed by the SRv6-VPN SID associated with the route. The underlay nodes whose SRv6 SID's are part of the SRH must support SRv6 data plane.

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

This document describes how existing BGP messages between PEs may carry SRv6 Segment IDs (SIDs) as a means to interconnect PEs and form VPNs.

## 2. SRv6-VPN SID TLV

The SRv6-VPN SID TLV is defined as another TLV for BGP-Prefix-SID Attribute [I-D.ietf-idr-bgp-prefix-sid]. The value field of the BGP Prefix SID attribute is defined here to be a set of elements encoded

as "Type/Length/Value" (i.e., a set of TLVs).  Type for SRv6-VPN SID TLV is defined to be TBD.

The IPv6-SID TLV MUST be present in the Prefix-SID attribute attached to MP-BGP VPN NLRI defined in [RFC4659][RFC5549][RFC7432] when egress-PE is capable of SRv6 data-plane.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       Type    |             Length            |   RESERVED    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  SRv6 SID information(Variable)                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

SRv6 SID information is encoded as follows:

```
+---------------------------------------+
|  SID Type (1 Octet)                   |
+---------------------------------------+
|  SRv6 SID (16 octet)                  |
+---------------------------------------+
```

Where:

- o  Type is TBD

- o  Length: 16bit field.  The total length of the value portion of the TLV.

- o  RESERVED: 8 bit field.  SHOULD be 0 on transmission and MUST be ignored on reception.

Current Type of SID defined as:

- o  Type-1 - corresponds to the equivalent functionality provided by a VPN MPLS Label attribute when received with a route containing a MPLS label[RFC4364].  Some functions which MAY be encoded are End.DX4, End.DT4, End.DX6, End.DT6 etc.

- o  Type-2 - corresponds to the equivalent functionality provided by a MPLS Label1 for EVPN Route-Types as defined in [RFC7432].  Some functions which MAY be encoded are End.DX2, End.DX2V, End.DT2U, End.DT2M / Arg.FE2 etc.

## 3. BGP based L3 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-PE) receive these advertisements and may add the prefix to the RIB in an appropriate VRF.

For egress-PEs which supports SRv6-VPN advertises an SRv6-VPN SID with VPN routes. This SRv6-VPN SID only has local significance at the egress-PE where it is allocated or configured on a per-CE or per-VRF basis. In practice, the SID encodes a cross-connect to a specific Address Family table (END.DT) or next-hop/interface (END.DX) as defined in the SRv6 Network Programming Document [I-D.filsfils-spring-srv6-network-programming]

The SRv6 VPN SID MAY be routable 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 VPN identifier.

To support SRv6 based L3VPN overlay, a SID is advertised with BGP MPLS L3VPN route update[RFC4364]. SID is encoded in a SRv6-VPN SID TLV, which is optional transitive BGP Prefix SID attribute[I-D.ietf-idr-bgp-prefix-sid]. This attribute serves two purposes; first it indicates that the BGP egress device is reachable via an SRv6 underlay and the BGP ingress device receiving this route MAY choose to encapsulate or insert an SRv6 SRH, second it indicates the value of the SID to include in the SRH encapsulation. For L3VPN, only a single SRv6-VPN SID MAY be necessary. A BGP speaker supporting an SRv6 underlay MAY distribute SID per route via the BGP SRv6-VPN Attribute. If the BGP speaker supports MPLS based L3VPN simultaneously, it MAY also populate the Label values in L3VPN route types and allow the BGP ingress device to decide which encapsulation to use. If the BGP speaker does not support MPLS based L3VPN services the MPLS Labels in L3VPN route types MUST be set to IMPLICIT-NULL. Similarly, to support SRv6 based EVPN a SID (or multiple SIDs) are advertised in route-types 1, 2, 3 and 5[RFC7432]

At an ingress-PE, BGP installs the advertised prefix in the correct RIB table, recursive via an SR Policy leveraging the received SRv6-VPN SID.

Assuming best-effort connectivity to the egress PE, the SR policy has a path with a SID list made up of a single SID: the SRv6-VPN SID received with the related BGP route update.

However, when VPN route is colored with an extended color community C and signaled with Next-Hop N and the ingress PE has a valid SRv6

Policy (N, C) associated with SID list <S1,S2, S3> [I-D.filsfils-spring-segment-routing-policy] then the SR Policy is <S1, S2, S3, SRv6-VPN SID>.

Multiple VPN routes MAY resolve recursively on the same SR Policy.

### 3.1. IPv4 VPN Over SRv6 Core

IPv4 VPN Over IPv6 Core is defined in [RFC5549], the MP_REACH_NLRI is encoded as follows for an SRv6 Core:

- o AFI = 1
- o SAFI = 128
- o Length of Next Hop Network Address = 16 (or 32)
- o Network Address of Next Hop = IPv6 address of the egress PE
- o NLRI = IPv4-VPN routes
- o Label = Implicit-Null

SRv6-VPN SID is encoded as part of the SRv6-VPN SID TLV defined in Section 2. The function of the SRv6 SID is entirely up to the originator of the advertisement. In practice, the function may likely be End.DX4 or End.DT4.

### 3.2. IPv6 VPN Over SRv6 Core

IPv6 VPN over IPv6 Core is defined in [RFC4659], the MP_REACH_NLRI is enclosed as follows for an SRv6 Core:

- o AFI = 2
- o SAFI = 128
- o Length of Next Hop Network Address = 16 (or 32)
- o Network Address of Next Hop = IPv6 address of the egress PE
- o NLRI = IPv6-VPN routes
- o Label = Implicit-Null

SRv6-VPN SID are encoded as part of the SRv6-VPN SID TLV defined in Section 2. The function of the IPv6 SRv6 SID is entirely up to the originator of the advertisement. In practice the function may likely be End.DX6 or End.DT6.

### 3.3. Global IPv4 over SRv6 Core

IPv4 over IPv6 Core is defined in [RFC5549]. The MP_REACH_NLRI is encoded with:

- AFI = 1
- SAFI = 1
- Length of Next Hop Network Address = 16 (or 32)
- Network Address of Next Hop = IPv6 address of Next Hop
- NLRI = IPv4 routes

SRv6 Global SID are encoded as part of the SRv6 SID TLV defined in Section 2. The function of the IPv6 SRv6 SID is entirely up to the originator of the advertisement. In practice, the function may likely be End.DX6 or End.DT6.

### 3.4. Global IPv6 over SRv6 Core

The MP_REACH_NLRI is encoded with:

- AFI = 2
- SAFI = 1
- Length of Next Hop Network Address = 16 (or 32)
- Network Address of Next Hop = IPv6 address of Next Hop
- NLRI = IPv6 routes

SRv6 Global SID are encoded as part of the SRv6 SID TLV defined in Section 2. The function of the IPv6 SRv6 SID is entirely up to the originator of the advertisement. In practice, the function may likely be End.DX6 or End.DT6.

Also, by utilizing the SRv6 SID TLV, as defined in Section 2, to encode the Global SID, BGP free core is possible by encapsulating all BGP traffic from edge to edge over SRv6.

## 4.  BGP based Ethernet VPN(EVPN) over SRv6

Ethernet VPN(EVPN), as defined in [RFC7432] provides an extendable method of building an EVPN overlay.  It primarily focuses on MPLS based EVPNs but calls out the extensibility to IP based EVPN overlays.  It defines 4 route-types which carry prefixes and MPLS Label attributes, the Labels each have specific use for MPLS encapsulation of EVPN traffic.  The fifth route-type carrying MPLS label information (and thus encapsulation information) for EVPN is defined in[I-D.ietf-bess-evpn-prefix-advertisement].  The Route Types discussed below are:

- Ethernet Auto-discovery Route

- MAC/IP Advertisement Route

- Inclusive Multicast Ethernet Tag Route

- Ethernet Segment route

- IP prefix route

- Selective Multicast route

- IGMP join sync route

- IGMP leave sync route

To support SRv6 based EVPN overlays a SID is advertised in route-type 1,2,3 and 5 above.  The SID (or SIDs) per route-type are advertised in a new SRv6-VPN SID TLV which is optional transitive BGP Prefix SID attribute.  This attribute serves two purposes; first it indicates that the BGP egress device is reachable via an SRv6 underlay and the BGP ingress device receiving this route MAY choose to encapsulate or insert an SRv6 SRH, second it indicates the value of the SID or SIDs to include in the SRH encapsulation.  A BGP speaker supporting an SRv6 underlay MAY distribute SIDs per route via the BGP SRv6 Attribute.  If the BGP speaker supports MPLS based EVPN simultaneously it MAY also populate the Label values in EVPN route types and allow the BGP ingress device to decide which encapsulation to use.  If the BGP speaker does not support MPLS based EVPN services the MPLS Labels in EVPN route types MUST be set to IMPLICIT-NULL.

### 4.1.  Ethernet Auto-discovery Route over SRv6 Core

Ethernet Auto-discovery (A-D) routes are Type-1 route type 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.

Multi-homed PEs MAY advertise an Ethernet auto discovery route per Ethernet segment with the introduced ESI MPLS label extended community defined in [RFC7432]. PEs may identify other PEs connected to the same Ethernet segment after the EVPN type-4 ES route exchange. All the multi-homed and remote PEs that are part of same EVI may import the auto discovery route.

EVPN Route Type-1 is encoded as follows for SRv6 Core:

```
+---------------------------------------+
|  RD (8 octets)                        |
+---------------------------------------+
|Ethernet Segment Identifier (10 octets)|
+---------------------------------------+
|  Ethernet Tag ID (4 octets)           |
+---------------------------------------+
|  MPLS label (3 octets)                |
+---------------------------------------+
```

For a SRv6 only BGP speaker for an SRv6 Core:

o  SRv6-VPN SID TLV MAY be advertised with the route.

#### 4.1.1. EVPN Route Type-1(Per ES AD)

Where:

o  BGP next-hop: IPv6 address of an egress PE

o  Ethernet Tag ID: all FFFF's

o  MPLS Label: always set to zero value

o  Extended Community: Per ES AD, ESI label extended community

SRv6-VPN TLV MAY be advertised along with the route advertisement and the behavior of the SRv6-VPN SID is entirely up to the originator of the advertisement. In practice, the behavior would likely be Arg.FE2.

#### 4.1.2. Prefix Type-1(Per EVI/ES AD)

Where:

o  BGP next-hop: IPv6 address of an egress PE

- o Ethernet Tag ID: non-zero for VLAN aware bridging, EVPN VPWS and FXC

- o MPLS Label: Implicit-Null

SRv6-VPN TLV MAY be advertised along with the route advertisement and the behavior of the SRv6-VPN SID is entirely up to the originator of the advertisement. In practice, the behavior would likely be END.DX2, END.DX2V or END.DT2U.

## 4.2. MAC/IP Advertisement Route(Type-2) with 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.

A MAC/IP Advertisement route type is encoded as follows for SRv6 Core:

```
+---------------------------------------+
|  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)          |
+---------------------------------------+
```

where:

- o BGP next-hop: IPv6 address of an egress PE

- o MPLS Label1: Implicit-null

- o MPLS Label2: Implicit-null

SRv6-VPN SID TLV MAY be advertised. The behavior of the SRv6-VPN SID is entirely up to the originator of the advertisement. In practice, the behavior of the SRv6 SID is as follows:

- END.DX2, END.DT2U (Layer 2 portion of the route)

- END.DT6/4 or END.DX6/4 (Layer 3 portion of the route)

Described below are different types of Type-2 advertisements.

- MAC/IP Advertisement Route(Type-2) with MAC Only

  * BGP next-hop: IPv6 address of egress PE

  * MPLS Label1: Implicit-null

  * MPLS Label2: Implicit-null

  * SRv6-VPN SID TLV MAY encode END.DX2 or END.DT2U behavior

- MAC/IP Advertisement Route(Type-2) with MAC+IP

  * BGP next-hop: IPv6 address of egress PE

  * MPLS Label1: Implicit-Null

  * MPLS Label2: Implicit-Null

  * SRv6-VPN SID TLV MAY encode Layer2 END.DX2 or END.DT2U behavior and Layer3 END.DT6/4 or END.DX6/4 behavior

### 4.3. Inclusive Multicast Ethernet Tag Route with 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.

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

An Inclusive Multicast Ethernet Tag route type specific EVPN NLRI consists of the following [RFC7432] where:

- o  BGP next-hop: IPv6 address of egress PE

- o  SRv6-VPN TLV MAY encode END.DX2/END.DT2M function.

- o  BGP Attribute: PMSI Tunnel Attribute[RFC6514] MAY contain MPLS implicit-null label and Tunnel Type would be similar to defined in EVPN Type-6 i.e. Ingress replication route.

The format of PMSI Tunnel Attribute attribute is encoded as follows for an SRv6 Core:

```
+--------------------------------------+
|  Flag (1 octet)                      |
+--------------------------------------+
|  Tunnel Type (1 octet)               |
+--------------------------------------+
|  MPLS label (3 octet)                |
+--------------------------------------+
|  Tunnel Identifier (variable)        |
+--------------------------------------+
```

- o  Flag: zero value defined per [RFC7432]

- o  Tunnel Type: defined per [RFC6514]

- o  MPLS label: Implicit-Null

- o  Tunnel Identifier: IP address of egress PE

SRv6 SID MAY be encoded as part of the SRv6-VPN SID TLV.  The behavior of the SRv6-VPN SID is entirely up to the originator of the advertisement.  In practice, the behavior of the SRv6 SID is as follows:

- o  END.DX2 or END.DT2M function

- o  The lower 32 bits of the SRv6-VPN SID TLV MAY be all zero's.  The ESI Filtering argument(Arg.FE2) carried along with EVPN Route Type-1 MAY be merged together by doing a bitwise logical OR to create a single SID on the ingress PE for Split-horizon and other filtering mechanisms.  Details of filtering mechanisms are described in[RFC7432]

### 4.4. Ethernet Segment Route with SRv6 Core

An Ethernet Segment route type specific EVPN NLRI consists of the following defined in [RFC7432]

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

where:

o  BGP next-hop: IPv6 address of egress PE

As oppose as previous route types, SRv6-VPN TLV is NOT advertised along with the route.  The processing of that route has not changed; it remains as described in [RFC7432].

### 4.5. IP prefix router(Type-5) with 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.  IP address may include host IP prefix or any specific subnet.  EVPN route Type-5 is defined in[I-D.ietf-bess-evpn-prefix-advertisement]

An IP Prefix advertisement is encoded as follows for an SRv6 Core:

```
+---------------------------------------+
|  RD (8 octets)                        |
+---------------------------------------+
|Ethernet Segment Identifier (10 octets)|
+---------------------------------------+
|  Ethernet Tag ID (4 octets)           |
+---------------------------------------+
|  IP Prefix Length (1 octet)           |
+---------------------------------------+
|  IP Prefix (4 or 16 octets)           |
+---------------------------------------+
|  GW IP Address (4 or 16 octets)       |
+---------------------------------------+
|  MPLS Label (3 octets)                |
+---------------------------------------+
```

- BGP next-hop: IPv6 address of egress PE

- MPLS Label: Implicit-Null

SRv6-VPN SID TLV MAY be advertised. The behavior of the SRv6-VPN SID is entirely up to the originator of the advertisement. In practice, the behavior of the SRv6 SID is an End.DT6/4 or End.DX6/4.

### 4.6. Multicast routes (EVPN Route Type-6, Type-7, Type-8)

These routes do not require any additional SRv6-VPN TLV. As per EVPN route-type 4, the BGP nexthop is equal to the IPv6 address of egress PE. More details may be added in future revisions of this document.

## 5. Migration from L3 MPLS based Segment Routing to SRv6 Segment Routing

Migration from MPLS to an SRv6 with BGP speakers is achieved with BGP sessions per BGP instance, one for IPv4 and a one for IPv6. Migration from IPv4 to IPv6 is independent of SRv6 BGP endpoints, and the selection of which route to use (received via the IPv4 or IPv6 session) is a local configurable decision of the ingress-PE, and is outside the scope of this document.

Migration from IPv6 MPLS based underlay to an SRv6 underlay with BGP speakers is achieved with a few simple rules at each BGP speaker.

```
 At Egress-PE
   If BGP offers an SRv6-VPN service
       Then BGP allocates an SRv6-VPN SID for the VPN service
       and adds the BGP SRv6-VPN SID TLV while advertising VPN prefixes.
   If BGP offers an MPLS VPN service
       Then BGP allocates an MPLS Label for the VPN service and
       use it in NLRI as normal for MPLS L3 VPNs.
   else MPLS label for VPN service is set to IMPLICIT-NULL.

 At Ingress-PE
   *Selection of which encapsulation below (SRv6-VPN or MPLS-VPN) is
    defined by local BGP policy
   If BGP supports SRv6-VPN service, and
   receives a BGP SRv6-VPN SID Attribute with an SRv6 SID
       Then BGP programs the destination prefix in RIB recursive via
       the related SR Policy.
   If BGP supports MPLS VPN service, and
   the MPLS Label is not Implicit-Null
       Then the MPLS label is used as a VPN label and inserted with the
       prefix into RIB via the BGP Nexthop.
```

## 6. Implementation Status

The SRv6-VPN is available for SRv6 on various Cisco hardware and other software platforms. An end-to-end integration of SRv6 L3VPN, SRv6 Traffic-Engineering and Service Chaining. All of that with data-plane interoperability across different implementations [1]:

o Three Cisco Hardware-forwarding platforms: ASR 1K, ASR 9k and NCS 5500

o Two Cisco network operating systems: IOS XE and IOS XR

o Barefoot Networks Tofino on OCP Wedge-100BF

o Linux Kernel officially upstreamed in 4.10

o Fd.io

## 7. Error Handling of BGP SRv6 SID Updates

When a BGP Speaker receives a BGP Update message containing a malformed SRv6-VPN SID TLV, it MUST ignore the received BGP attributes and not pass it to other BGP peers. This is equivalent to the -attribute discard- action specified in [RFC7606]. When discarding an attribute, a BGP speaker MAY log an error for further analysis.

## 8. IANA Considerations

This document defines a new TLV types as part of the BGP Prefix SID attribute.

## 9. Security Considerations

This document introduces no new security considerations beyond those already specified in [RFC4271] and [RFC3107].

## 10. Conclusions

This document proposes extensions to the BGP to allow advertising certain attributes and functionalities related to SRv6.

## 11. References

## 11.1. Normative References

[I-D.filsfils-spring-segment-routing-policy]
Filsfils, C., Sivabalan, S., Raza, K., Liste, J., Clad, F., Talaulikar, K., Hegde, S., daniel.voyer@bell.ca, d., Lin, S., bogdanov@google.com, b., Horneffer, M., Steinberg, D., Decraene, B., Litkowski, S., and P. Mattes, "Segment Routing Policy for Traffic Engineering", draft-filsfils-spring-segment-routing-policy-04 (work in progress), December 2017.

[I-D.filsfils-spring-srv6-network-programming]
Filsfils, C., Leddy, J., daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d., Steinberg, D., Raszuk, R., Matsushima, S., Lebrun, D., Decraene, B., Peirens, B., Salsano, S., Naik, G., Elmalky, H., Jonnalagadda, P., Sharif, M., Ayyangar, A., Mynam, S., Henderickx, W., Bashandy, A., Raza, K., Dukes, D., Clad, F., and P. Camarillo, "SRv6 Network Programming", draft-filsfils-spring-srv6-network-programming-03 (work in progress), December 2017.

[I-D.ietf-6man-segment-routing-header]
Previdi, S., Filsfils, C., Raza, K., Leddy, J., Field, B., daniel.voyer@bell.ca, d., daniel.bernier@bell.ca, d., Matsushima, S., Leung, I., Linkova, J., Aries, E., Kosugi, T., Vyncke, E., Lebrun, D., Steinberg, D., and R. Raszuk, "IPv6 Segment Routing Header (SRH)", draft-ietf-6man-segment-routing-header-07 (work in progress), July 2017.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <https://www.rfc-editor.org/info/rfc2460>.

[RFC3107] Rekhter, Y. and E. Rosen, "Carrying Label Information in BGP-4", RFC 3107, DOI 10.17487/RFC3107, May 2001, <https://www.rfc-editor.org/info/rfc3107>.

[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>.

[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>.

### 11.2. Informative References

[I-D.ietf-bess-evpn-prefix-advertisement]
Rabadan, J., Henderickx, W., Drake, J., Lin, W., and A. Sajassi, "IP Prefix Advertisement in EVPN", draft-ietf-bess-evpn-prefix-advertisement-09 (work in progress), November 2017.

[I-D.ietf-idr-bgp-prefix-sid]
Previdi, S., Filsfils, C., Lindem, A., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix SID extensions for BGP", draft-ietf-idr-bgp-prefix-sid-07 (work in progress), October 2017.

[I-D.ietf-isis-segment-routing-extensions]
Previdi, S., Ginsberg, L., Filsfils, C., Bashandy, A., Gredler, H., Litkowski, S., Decraene, B., and J. Tantsura, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-15 (work in progress), December 2017.

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", draft-ietf-spring-segment-routing-14 (work in progress), December 2017.

[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>.

[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>.

[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>.

[RFC5549] Le Faucheur, F. and E. Rosen, "Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop", RFC 5549, DOI 10.17487/RFC5549, May 2009, <https://www.rfc-editor.org/info/rfc5549>.

## 11.3. URIs

[1] http://www.segment-routing.net

## Appendix A. Contributors

Bart Peirens
Proximus
Belgium

Email: bart.peirens@proximus.com

Authors' Addresses

Gaurav Dawra (editor)
Cisco Systems
USA

Email: gdawra.ietf@gmail.com

Clarence Filsfils
Cisco Systems
Belgium

Email: cfilsfil@cisco.com

Darren Dukes
Cisco Systems
Canada

Email: ddukes@cisco.com

Patrice Brissette
Cisco Systems
Canada

Email: pbrisset@cisco.com

Pablo Camarilo
Cisco Systems
Spain

Email: pcamaril@cisco.com

Jonn Leddy
Comcast
USA

Email: john_leddy@cable.comcast.com

Daniel Voyer
Bell Canada
Canada

Email: daniel.voyer@bell.ca

Daniel Bernier
Bell Canada
Canada

Email: daniel.bernier@bell.ca

Dirk Steinberg
Steinberg Consulting
Germany

Email: dws@steinberg.net

Robert Raszuk
Bloomberg LP
USA

Email: robert@raszuk.net

Bruno Decraene
Orange
France

Email: bruno.decraene@orange.com

Satoru Matsushima
SoftBank
1-9-1,Higashi-Shimbashi,Minato-Ku
Japan 105-7322

Email: satoru.matsushima@g.softbank.co.jp