

IDR
Internet-Draft
Intended status: Standards Track
Expires: March 11, 2021

S. Sangli
R. Bonica
Juniper Networks Inc.
September 07, 2020

BGP based Virtual Private Network (VPN) Services over SRm6 enabled IPv6
networks

[draft-ssangli-bess-bgp-vpn-srm6-02](#)

Abstract

This document defines BGP protocol extensions for encoding and carrying SRm6 Tunnel Payload Forwarding information (TPF) to support Virtual Private Network services. This is applicable when the VPN services are offered in a SRm6 enabled IPv6 network such that the VPN payload is transported over IPv6. The Tunnel Payload Information is encoded in the IPv6 Destination Option Header in the IPv6 data packets.

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 March 11, 2021.

Copyright Notice

Copyright (c) 2020 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	2
2.	Requirements Language	3
3.	Per-Path Service Instruction Information	3
4.	Usage of Tunnel Encapsulation Attribute	4
5.	Procedures for Egress BGP Speaker	6
6.	Procedures for Ingress BGP Speaker	7
7.	BGP Nexthop and Tunnel Endpoint address handling procedures	7
8.	BGP based L3 VPN services over IPv6	8
8.1.	IPv4 VPN on SRm6 enabled IPv6 Core	8
8.2.	IPv6 VPN on SRm6 enabled IPv6 Core	8
8.3.	IPv4 Global Routes on SRm6 enabled IPv6 Core	9
9.	BGP based Ethernet VPN services over IPv6	9
9.1.	Ethernet Per ES Auto-Discovery (A-D) route	10
9.2.	Ethernet per EVI Auto-Discovery (A-D) route	10
9.3.	MAC/IP Advertisement route	11
9.4.	Inclusive Multicast Ethernet Route	11
9.5.	IP Prefix Route	12
10.	Deployment Considerations	12
11.	Backward Compatibility	13
12.	Security Considerations	14
13.	IANA Considerations	14
14.	Acknowledgements	14
15.	References	14
15.1.	Normative References	14
15.2.	Informative References	15
	Authors' Addresses	17

[1. Introduction](#)

Virtual Private Network (VPN) technologies allow network providers to emulate private networks with shared infrastructure. For example, assume that a set of red sites, set of blue sites and a set of green sites connect to a provider network. Furthermore, assume that red sites and blue sites wish to interconnect, exchange packets.

However, the green sites wish to communicate with green sites only. The provider should allow its infrastructure network to scale to both the requirements without having to create multiple parallel network infrastructures. The IETF has standardized many VPN technologies viz. Layer 3 VPN (L3VPN) [[RFC4364](#)], Layer 2 VPN (L2VPN) [[RFC6624](#)], Virtual Private LAN Service (VPLS) [[RFC4761](#)], [[RFC4762](#)], Ethernet VPN

(EVPN) [[RFC7432](#)], Pseudowires [[RFC8077](#)] to enable Layer 3 and Layer 2 VPN services.

The aforementioned technologies leverage MPLS network architecture :

- o to establish a MPLS tunnel from ingress PE to egress PE, thus making all P routers agnostic of VPN state.
- o to provide demultiplexing abstraction in the tunnelled packet so the payload packet can be forwarded at the egress router based on Routing table and/or interface.

In pure IPv6 deployments where there may be non-MPLS capable routers, it would be desirable to have alternate mechanism to provide VPN connectivity. This document describes BGP extensions and procedures applicable for SRm6 enabled IPv6 networks, to provide VPN services over BGP.

2. 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.

3. Per-Path Service Instruction Information

A SRm6 [[I-D.bonica-spring-sr-mapped-six](#)] segment provides unidirectional connectivity from an ingress node to an egress node. A SRm6 path contains one or more such segments. SRm6 introduces the concept of Per-Segment Service Instruction and Per-Path Service Instruction. These instructions describe the additional packet processing performed on a node. The Per-Segment Service Instruction is executed on the segment egress node while the Per-Path Service Instruction is executed on the path egress node. The SR Path egress node advertises the service prefix reachability information to SR Path ingress node via Multi-Protocol extensions in BGP [[RFC4760](#)].

For providing VPN services, aforementioned BGP extensions rely on MPLS architecture [[RFC3031](#)]. The BGP extensions specify the new encoding for Network Layer Reachability Information (NLRI) to include the MPLS VPN labels [[RFC8277](#)]. Such a MPLS VPN label is associated with a forwarding decision in the VPN Routing Instance on the egress BGP Router. The ingress BGP router will push the VPN label on the data packet destined to the egress BGP router. The transport tunnel from ingress router to egress router can be MPLS or GRE or L2TPv3, but inner payload is a MPLS packet as described in [[RFC4023](#)],

[RFC4817], [RFC7510]. The intermediate routers do not process the VPN label [a.k.a.] embedded label as described in [I-D.ietf-idr-tunnel-encaps].

To provide BGP based VPN services on a non-MPLS IPv6 networks, it would be beneficial to retain the benefits of BGP protocol extensions while leveraging the benefits of IPv6 [RFC8200].

[I-D.bonica-6man-vpn-dest-opt] describes SRm6 paths as programmable with Tunnel Payload Forwarding information (TPF) that determine how egress nodes process SRm6 payloads. The TPF information is carried in the Tunnel Payload Forwarding Option encoded in the IPv6 Destination Option Header [RFC8200].

The Tunnel Payload Forwarding (TPF) information is defined as follows:

- o 32 bit quantity.

The TPF information have node-local significance and is assigned by the egress BGP router. The value of zero is reserved. The TPF information will serve 2 purposes.

- o It MUST uniquely identify the VPN Routing Instance for L3VPN or identify an Ethernet Segment for EVPN or identify a leaf property for EVPN TREE upon which forwarding decision can be taken.
- o It MAY provide information for special processing before the packet is forwarded.

The structure of TPF information will be updated in the next version of this document.

The encoding of the Tunnel Payload Forwarding information for VPNs is described in [Section 8](#) and [Section 9](#).

4. Usage of Tunnel Encapsulation Attribute

This document defines a new Tunnel type : SRm6. The format is as per below.

- o Tunnel Type (2 Octets) : To be assigned
- o Tunnel Length (2 Octets) : 1
- o Value : List of Sub-TLVs

[I-D.ietf-idr-tunnel-encaps] defines many sub-TLVs for the tunnels. The encoding for them are as follows:

- o Tunnel Endpoint Sub-TLV : As per [[I-D.ietf-idr-tunnel-encaps](#)]
- o Encapsulation Sub-TLV : Not needed.
- o IPv4 DS Field Sub-TLV : Not needed.
- o UDP Destination Port Sub-TLV : Not needed.
- o Protocol Type Sub-TLV : As per [[I-D.ietf-idr-tunnel-encaps](#)].
- o Color Sub-TLV : As per [[I-D.ietf-idr-tunnel-encaps](#)].
- o Embedded Label Handling Sub-TLV : 3.
- o MPLS Label Stack Sub-TLV : Not needed.
- o Prefix SID Sub-TLV : Not Needed.

The Tunnel Encapsulation Attribute is a an Optional Transitive attribute as described in [[I-D.ietf-idr-tunnel-encaps](#)]. This attribute with SRm6 tunnel type MUST be present in the BGP update carrying the Network Layer Reachability Information encoded with the TPF information. This document refers to the NLRI that is associated with SRm6 Tunnel Encapsulation attribute as SRm6_NLRI. The document [[I-D.ietf-idr-tunnel-encaps](#)] defines the encoding for sub-TLV as follows.

- o Sub-TLV Type : 1 octet
- o Sub-TLV Length : 1 or 2 octets
- o Sub-TLV Value : defined per Sub-TLV as per below.

The Tunnel Endpoint Sub-TLV can specify the IPv6 address of the egress router as the final destination address of SRm6 packet which is also referred to as SR Path destination address. The sub-fields on this sub-TLV is encoded as below.

- o Autonomous System Number : AS number of the IPv6 SR domain.
- o Address Family : 2 (refers to IPv6).
- o Address : IPv6 address of the egress interface present in SRm6 domain.

The Value field may be set to 0 which indicates that next hop value in the NLRI should be chosen for the SRm6 Path destination address.

The Embedded Label Handling Sub-TLV describes how the label field in the NLRI should be interpreted.

- o Value : MUST be set to 3.

The [[I-D.ietf-idr-tunnel-encaps](#)] specifies only 2 values. While the value 1 refers to label field as MPLS embedded label that is carried at the top of the label stack of the MPLS payload packet, the value 2 refers to label field to be either ignored or carried in the virtual network field of the encapsulation header.

This document defines another behavior for the label field. The value 3 will indicate that value in the label field MUST be inserted in the Destination Options Header of the IPv6 Tunnel header.

The Tunnel Encapsulation attribute can carry one or more Tunnel types. The local policy on the ingress router can determine which Tunnel type to be used for the NLRI. The Tunnel Endpoint address MUST be set only by the egress BGP router that is the endpoint of the SRm6 path.

5. Procedures for Egress BGP Speaker

The TPF information instructs the egress router to de-encapsulate the packet and forward the newly exposed payload inner packet through the specified interface or forward using the specified Routing Instance. The TPF information described in [Section 3](#) will be assigned by the egress BGP Router.

When the egress BGP Speaker advertises the NLRI, it will include the TPF information in the encoding described in [Section 8](#) and [Section 9](#). The egress BGP Speaker MUST include the Tunnel Encapsulation Attribute with Route type SRm6 as described in [Section 4](#) in such BGP updates.

By tagging the BGP update with Tunnel Encapsulation attribute of SRm6 type, the BGP Speaker informs how the SRm6_NLRI should be decoded and processed by the receiving BGP Speaker.

Via the Remote Tunnel Endpoint Sub-TLV encoding, the egress BGP router may specify the SRm6 Path Destination Address. The Protocol type Sub-TLV and the Color Sub-TLV may be used by the egress BGP router to influence the payload packets to be put on SRm6 path. The Embedded Label Handling Sub-TLV MUST be set to 3 to inform that the label field MUST be used to form the TPF option that is inserted in the Destination Options Header at the ingress router as described in [[I-D.bonica-6man-vpn-dest-opt](#)].

A single TPF information may be associated with all the prefixes in a Routing Instance or a unique TPF information may be associated for each prefix in the Routing Instance. Similarly, a TPF information may be assigned to identify an Ethernet segment or leaf AC property by EVPN. The choice is left to the Network Operator and is outside the scope of this document.

6. Procedures for Ingress BGP Speaker

Upon receiving a BGP update, the receiving BGP Speaker will look for Tunnel Encapsulation attribute. If the tunnel type carried in the Tunnel Encapsulation attribute is SRm6, the BGP updates is said to be carrying the SRm6_NLRI and the Label field in the Network Layer Reachability Information is treated as Tunnel Payload Forwarding information (TPF).

The tuple (TPF information, Prefix) is programmed in the forwarding infrastructure of the router. The manner in which this tuple is stored in the router is outside the scope of this document. If SRm6 has been enabled on the router, such a tuple SHOULD be used for encoding the Destination Options Header as described in [\[I-D.bonica-6man-vpn-dest-opt\]](#).

The [\[I-D.ietf-idr-tunnel-encaps\]](#) describes how Tunnel Endpoint Sub-TLV has to be processed. It also describes the usage of the Protocol type Sub-TLV and the Color Sub-TLV. This may be used by the ingress BGP router to select the payload packets that should be put on SRm6 path.

The Embedded Label Handling Sub-TLV value that is set to 3 indicates that ingress BGP router to use the value of label field to construct the Tunnel Payload Forwarding Option that is inserted in the Destination Options Header of the Tunnel IPv6 packet.

7. BGP Nexthop and Tunnel Endpoint address handling procedures

The BGP Nexthop attribute handling procedures are described in [\[RFC4271\]](#) while [\[RFC4760\]](#) describe the handling procedures for the Nexthop field in the MP_REACH attribute. The target="I-D.ietf-idr-tunnel-encaps"/> describes the Tunnel Endpoint sub-TLV in the Tunnel Encapsulation Attribute as the next hop address to which the prefix should be forwarded to. If a BGP update has such a Tunnel Encapsulation Attribute it prescribes that the Tunnel Endpoint Sub-TLV if non-zero, MUST be used as the next hop to send the packet to.

There may be instances where the BGP update carrying the SRm6 NLRI will cross Autonomous boundary. The BGP update with SRm6 NLRI MUST always carry the Tunnel Encapsulation Attribute. If any router along

the path wishes to change the Tunnel Endpoint Sub-TLV next hop address, it MUST also update the TPF information field of the The BGP update carrying the SRm6 NLRI.

It should be noted that router that modifies the Tunnel Endpoint sub-TLV of the Tunnel Encapsulation attribute present in the SRm6 update must be able to stitch the egress tunnel and ingress tunnel.

8. BGP based L3 VPN services over IPv6

The Egress and Ingress BGP speakers form a BGP peering session to exchange a set of prefixes described in [[RFC4271](#)] and Multi-Protocol extensions [[RFC4760](#)]. The BGP Router capable of SRm6 that is enabled to carry L3 VPN services over IPv6 networks should follow the procedures mentioned in [Section 5](#) and [Section 6](#). The manner in which a BGP Router is configured for SRm6 underlay and L3 VPN overlay is outside the scope of this document.

8.1. IPv4 VPN on SRm6 enabled IPv6 Core

The IPv4 L3 VPN over IPv6 is defined in [[RFC5549](#)]. The MP_REACH NLRI and Tunnel Encapsulation attribute encoding is as per below:

- o AFI : 1; SAFI : 128
- o Length of the Next Hop : 16 (or 32 if Link Local)
- o Network address of the Next Hop : IPv6 address of the egress BGP Router
- o NLRI : IPv4-VPN routes
- o Label : Low order 24 bits of Tunnel Payload Forwarding (TPF) information
- o Tunnel Encapsulation Path Attribute : SRm6 Type as described in [Section 4](#)

The TPF information is associated with VPN Routing Instance on the Egress PE. The Tunnel Encapsulation attribute with SRm6 type MUST be appended to the Path attributes associated with the NLRI.

8.2. IPv6 VPN on SRm6 enabled IPv6 Core

The IPv6 L3 VPN over IPv6 is defined in [[RFC4659](#)]. The MP_REACH NLRI and Tunnel Encapsulation attribute encoding is as per below:

- o AFI : 2; SAFI : 128

- o Length of the Next Hop : 16 (or 32 if Link Local)
- o Network address of the Next Hop : IPv6 address of the egress BGP Router
- o NLRI : IPv6-VPN routes
- o Label : Low order 24 bits of Tunnel Payload Forwarding (TPF) information
- o Tunnel Encapsulation Path Attribute : SRm6 Type as described in ([Section 4](#))

The TPF information is associated with VPN Routing Instance on the Egress PE. The Tunnel Encapsulation attribute with SRm6 type MUST be appended to the Path attribute associated with the NLRI.

8.3. IPv4 Global Routes on SRm6 enabled IPv6 Core

The IPv4 L3 VPN over IPv6 is defined in [[RFC5549](#)]. The MP_REACH NLRI and Tunnel Encapsulation attribute encoding is per below:

- o AFI : 1; SAFI : 1
- o Length of the Next Hop : 16 (or 32 if Link Local)
- o Network address of the Next Hop : IPv6 address of the egress BGP Router
- o NLRI : IPv4 routes
- o Label : Low order 24 bits of Tunnel Payload Forwarding (TPF) information
- o Tunnel Encapsulation Path Attribute : SRm6 Type as described in ([Section 4](#))

The TPF information is associated with VPN Routing Instance on the Egress PE. The Tunnel Encapsulation attribute with SRm6 type MUST be appended to the Path attribute associated with the NLRI.

9. BGP based Ethernet VPN services over IPv6

The [[RFC7432](#)] describes the BGP extensions for carrying the Ethernet Virtual Private Network Overlay on MPLS network. It defines 4 types of EVPN NLRI. This document specifies changes to certain fields for those NLRIs.

- o Ethernet Auto-Discovery (A-D) route
- o MAC/IP Advertisement route
- o Inclusive Multicast Ethernet Tag route
- o IP Prefix route

9.1. Ethernet Per ES Auto-Discovery (A-D) route

The MP_REACH and MP_UNREACH attributes will carry this route in the NLRI encoding described in [RFC7432]. In addition to Tunnel Encapsulation attribute encoding, this document recommends to follow the [RFC7432] encoding except the following. For MPLS label carried in the Ethernet A-D per ESI route:

- o MPLS label : Per [RFC7432], it is set to zero.
- o Tunnel Encapsulation Path Attribute : SRm6 Type as described in (Section 4)

The MPLS label field is not part of the route but treated as route attribute. For procedures and usage of this route, refer to [RFC7432]. The Tunnel Encapsulation attribute with SRm6 type MUST be appended to the Path attribute associated with the NLRI.

An EVPN Ethernet per ES A-D route is usually signaled together with an ESI label extended community. For ESI Label carried in the ESI label extended community:

- o ESI Label: Low order 24 bits of the Tunnel Payload Forwarding (TPF) information

The TPF information is used to identify an Ethernet segment attached to the BGP PE for EVPN.

9.2. Ethernet per EVI Auto-Discovery (A-D) route

The MP_REACH and MP_UNREACH attributes will carry this route in the NLRI encoding described in [RFC7432]. In addition to Tunnel Encapsulation attribute encoding, this document recommends to follow the [RFC7432] encoding except the following:

- o MPLS label : Low order 24 bits of Tunnel Payload Forwarding (TPF) information
- o Tunnel Encapsulation Path Attribute : SRm6 Type as described in (Section 4)

The MPLS label field is not part of the route but treated as route attribute. For procedures and usage of this route, refer to [\[RFC7432\]](#). The Tunnel Encapsulation attribute with SRm6 type MUST be appended to the Path attribute associated with the NLRI.

In addition, for EVPN E-tree service, this route may be signaled together with an E-Tree Extended Community as it is specified in [\[RFC8317\]](#). For the leaf label carried in the E-Tree Extended Community:

- o Leaf Label: Low order 24 bits of the Tunnel Payload Forwarding (TPF) information

In case of EVPN E-tree service, the TPF information carried in the E-Tree extended community is used to signal a leaf AC property.

In the data plane, this TPF information specified in the Destination Option header is used by an egress router to identify that a data packet is ingressed from a leaf AC such that appropriate forwarding decision can be made.

9.3. MAC/IP Advertisement route

The MP_REACH and MP_UNREACH attributes will carry this route in the NLRI encoding described in [\[RFC7432\]](#). In addition to Tunnel Encapsulation attribute encoding, this document recommends to follow the [\[RFC7432\]](#) encoding except the following.

- o MPLS label1 : Low order 24 bits of the Tunnel Payload Forwarding (TPF) information1
- o MPLS label2 : Low order 24 bits of the Tunnel Payload Forwarding (TPF) information2
- o Tunnel Encapsulation Path Attribute : SRm6 Type as described in ([Section 4](#))

The MPLS label field is not part of the route but treated as route attribute. For procedures and usage of this route, refer to [\[RFC7432\]](#). The Tunnel Encapsulation attribute with SRm6 type MUST be appended to the Path attribute associated with the NLRI.

9.4. Inclusive Multicast Ethernet Route

The MP_REACH and MP_UNREACH attributes will carry this route in the NLRI encoding described in [\[RFC7432\]](#). In addition to Tunnel Encapsulation attribute encoding, this document recommends to follow the [\[RFC7432\]](#) encoding except the following.

- o If MPLS label field in the PMSI Tunnel Attribute is non-zero, it is set to Low order 24 bits of the Tunnel Payload Forwarding (TPF) information.
- o Tunnel Encapsulation Path Attribute : SRm6 Type as described in ([Section 4](#))

The Tunnel Encapsulation attribute with SRm6 type MUST be appended to the Path attribute associated with the NLRI.

9.5. IP Prefix Route

The MP_REACH and MP_UNREACH attributes will carry this route in the NLRI encoding described in [[I-D.ietf-bess-evpn-prefix-advertisement](#)]. In addition to Tunnel Encapsulation attribute encoding, this document recommends the following change:

- o MPLS label: if it is non-zero, it is set to Low order 24 bits of the Tunnel Payload Forwarding (TPF) information.
- o Tunnel Encapsulation Path Attribute : SRm6 Type as described in ([Section 4](#))

The MPLS label field is not part of the route but treated as route attribute. For procedures and usage of this route, refer to [[I-D.ietf-bess-evpn-prefix-advertisement](#)]. The Tunnel Encapsulation attribute with SRm6 type MUST be appended to the Path attribute associated with the NLRI.

10. Deployment Considerations

This document proposes to reuse the NLRI encoding for BGP L3VPN and EVPN Network Layer Routing Information. However, care should be taken when BGP VPN overlay services are enabled on SRm6 underlay such that Tunnel Encapsulation Path attribute with SRm6 type MUST be appended. When a BGP router advertises SRm6_NLRI, it MUST NOT remove the Tunnel Encapsulation Path attribute.

The SRm6 underlay is similar to other "tunnel" technologies viz MPLS, GRE, IP-in-IP, L2TPv3. The egress and ingress BGP routers can be connected via one or more such underlay technologies. A BGP speaker can advertise the VPN NLRI with the nexthop reachable via one or more such underlay paths. Each such mechanism can co-exist together as ships-in-night. However, when SRm6_NLRI is advertised by a egress BGP speaker and received by an ingress BGP speaker, they MUST follow the procedures mentioned in this document.

For migrating a BGP router to SRm6 the following procedures can be followed.

- o Operator will enable SRm6 underlay on the ingress and egress routers identifying the SRm6 path from ingress router's interface to egress router's interface. The way to configure the ingress and egress routers are outside the scope of this document.
- o SRm6 enabled ingress BGP router will setup the additional information in the forwarding table such that it can append an IPv6 tunnel header and encode the TPF Option in the Destination Options Header.
- o SRm6 enabled egress BGP router will setup the additional information in the forwarding table such that TPF information can be used to lookup to find the Routing Instance and make the forwarding decision.
- o Operator will enable BGP VPN overlay over SRm6 underlay on ingress router. This means that ingress router will start looking for SRm6_NLRI in the BGP updates. The way to enable the BGP VPN overlay over SRm6 underlay is outside the scope of this document.
- o The operator will enable BGP VPN overlay over SRm6 underlay on egress router. With this, the egress router will create TPF information and associate it with Routing Instances. It then advertises the SRm6_NLRIs to the ingress BGP router.
- o The ingress router will interpret the SRm6_NLRIs and use TPF information and follow the procedures in [\[I-D.bonica-spring-sr-mapped-six\]](#) to encode the Destination Options Header to forward the data packet.
- o Now that SRm6 path is setup between ingress and egress BGP routers, on the egress BGP router the Operator can migrate the Routing Instances from MPLS VPN set of Instances to SRm6 enabled set of Instances. The way to configure Routing Instances to achieve the above is outside the scope of this document.

11. Backward Compatibility

The extension proposed in this document is backward compatible with procedures described for BGP enabled services.

12. Security Considerations

This document does not introduce any new security considerations beyond those already specified in [[RFC4271](#)], [[RFC8277](#)] and [[I-D.ietf-idr-tunnel-encaps](#)].

13. IANA Considerations

IANA is requested to assign a code point for SRm6 Route Type for BGP Tunnel Encapsulation Path Attribute from BGP Tunnel Encapsulation Attribute Tunnel Types Registry.

14. Acknowledgements

The authors would like to thank Jeff Haas, Wen Lin and Shraddha Hegde for careful review and suggestions.

15. References

15.1. Normative References

- [I-D.bonica-6man-vpn-dest-opt]
Bonica, R., Kamite, Y., Jalil, L., Zhou, Y., and G. Chen,
"The IPv6 Tunnel Payload Forwarding (TPF) Option", [draft-bonica-6man-vpn-dest-opt-13](#) (work in progress), August 2020.
- [I-D.bonica-spring-sr-mapped-six]
Bonica, R., Hegde, S., Kamite, Y., Alston, A., Henriques, D., Jalil, L., Halpern, J., Linkova, J., and G. Chen,
"Segment Routing Mapped To IPv6 (SRm6)", [draft-bonica-spring-sr-mapped-six-01](#) (work in progress), April 2020.
- [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-11](#) (work in progress), May 2018.
- [I-D.ietf-idr-tunnel-encaps]
Patel, K., Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", [draft-ietf-idr-tunnel-encaps-17](#) (work in progress), July 2020.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.

- [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>>.
- [RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", [RFC 4303](#), DOI 10.17487/RFC4303, December 2005, <<https://www.rfc-editor.org/info/rfc4303>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, [RFC 4443](#), DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [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>>.

15.2. Informative References

- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", [RFC 3031](#), DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", [RFC 4023](#), DOI 10.17487/RFC4023, March 2005, <<https://www.rfc-editor.org/info/rfc4023>>.
- [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>>.

- [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>>.
- [RFC4761] Kompella, K., Ed. and Y. Rekhter, Ed., "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", [RFC 4761](#), DOI 10.17487/RFC4761, January 2007, <<https://www.rfc-editor.org/info/rfc4761>>.
- [RFC4762] Lasserre, M., Ed. and V. Kompella, Ed., "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", [RFC 4762](#), DOI 10.17487/RFC4762, January 2007, <<https://www.rfc-editor.org/info/rfc4762>>.
- [RFC4817] Townsley, M., Pignataro, C., Wainner, S., Seely, T., and J. Young, "Encapsulation of MPLS over Layer 2 Tunneling Protocol Version 3", [RFC 4817](#), DOI 10.17487/RFC4817, March 2007, <<https://www.rfc-editor.org/info/rfc4817>>.
- [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>>.
- [RFC6624] Kompella, K., Kothari, B., and R. Cherukuri, "Layer 2 Virtual Private Networks Using BGP for Auto-Discovery and Signaling", [RFC 6624](#), DOI 10.17487/RFC6624, May 2012, <<https://www.rfc-editor.org/info/rfc6624>>.
- [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>>.
- [RFC7510] Xu, X., Sheth, N., Yong, L., Callon, R., and D. Black, "Encapsulating MPLS in UDP", [RFC 7510](#), DOI 10.17487/RFC7510, April 2015, <<https://www.rfc-editor.org/info/rfc7510>>.

- [RFC8077] Martini, L., Ed. and G. Heron, Ed., "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", STD 84, [RFC 8077](#), DOI 10.17487/RFC8077, February 2017, <<https://www.rfc-editor.org/info/rfc8077>>.
- [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>>.

Authors' Addresses

Srihari Sangli
Juniper Networks Inc.
Exora Business Park
Bangalore, KA 560103
India

Email: ssangli@juniper.net

Ron Bonica
Juniper Networks Inc.
2251 Corporate Park Drive
Herndon, Virginia 20171
USA

Email: rbonica@juniper.net