IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 25, 2020

J. Dong
Z. Hu
Z. Li
Huawei Technologies
X. Tang
R. Pang
China Unicom
June 23, 2020

BGP-LS Extensions for Segment Routing based Enhanced VPN
draft-dong-idr-bgpls-sr-enhanced-vpn-02

Abstract

Enhanced VPN (VPN+) is an enhancement to VPN services to support the needs of new applications, particularly including the applications that are associated with 5G services. These applications require enhanced isolation and have more stringent performance requirements than that can be provided with traditional overlay VPNs. An enhanced VPN may be used for 5G transport network slicing, and will also be of use in more generic scenarios. To meet the requirement of enhanced VPN services, a number of Virtual Transport Networks (VTNs) need to be created, each with a subset of the underlay network topology and a set of network resources allocated to meet the requirement of a specific VPN+ service, or a group of VPN+ services.

This document specifies the BGP-LS mechanisms with necessary extensions to advertise the information of Segment Routing (SR) based VTNs. The proposed mechanism is applicable to both segment routing with MPLS data plane (SR-MPLS) and segment routing with IPv6 data plane (SRv6).

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Table of Contents

1. Introduction

Enhanced VPN (VPN+) is an enhancement to VPN services to support the needs of new applications, particularly including the applications that are associated with 5G services. These applications require enhanced isolation and have more stringent performance requirements than that can be provided with traditional overlay VPNs. These properties cannot be met with pure overlay networks, as they require integration between the underlay and the overlay networks. [I-D.ietf-teas-enhanced-vpn] specifies the framework of enhanced VPN and describes the candidate component technologies in different network planes and layers. An enhanced VPN can be used for 5G transport network slicing, and will also be of use in more generic scenarios.

To meet the requirement of enhanced VPN services, a number of Virtual Transport Networks (VTNs) need to be created, each with a subset of the underlay network topology and a set of network resources allocated to meet the requirement of a specific VPN+ service or a group of VPN+ services.

[I-D.dong-spring-sr-for-enhanced-vpn] specifies how segment routing (SR) [RFC8402] can be used to build virtual transport networks (VTNs) with the required network topology and network resources, which could be used as the underlay of enhanced VPN services. [I-D.dong-lsr-sr-enhanced-vpn] specifies the IGP mechanism and extensions to build a set of SR based VTNs. When a VTN spans multiple IGP areas or multiple Autonomous Systems (ASes), BGP-LS is needed to advertise the VTN information in each IGP area or AS to the network controller, so that the controller could use the collected information to build the inter-area or inter-AS SR VTNs.

This document describes BGP-LS [RFC7752] based mechanism with necessary extensions to advertise the topology and resource attribute of inter-area and inter-domain SR based VTNs. The definition of VTN is advertised as a node attribute using BGP-LS. The associated network topology and resources attributes of a VTN are advertised as link attribute using BGP-LS.

2. Advertisement of VTN Definition

According to [I-D.ietf-teas-enhanced-vpn], a VTN has a customized network topology and a set of dedicated or shared network resources. Thus a VTN can be defined as the combination of a set of network attributes, which include the topology attribute and other attributes, such as the associated network resources.

Virtual Transport Network Definition (VTND) TLV is a new TLV of the optional BGP-LS Attribute which is associated with the node NLRI.

The format of VTND TLV is as follows:

```
 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            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             VTN ID                            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             MT-ID             |   Algorithm   |     Flags     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Sub-TLVs                            |
~                              ...                              ~
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Where:

- Type: TBD

- Length: the length of the value field of the TLV. It is variable dependent on the included Sub-TLVs.

- VTN ID: A global significant 32-bit identifier which is used to identify a virtual transport network.

- MT-ID: 16-bit identifier which indicates the multi-topology identifier of the IGP topology.

- Algorithm: 8-bit identifier which indicates the algorithm which applies to this virtual transport network. It can be either a normal algorithm in [RFC8402] or a Flex-Algorithm [I-D.ietf-lsr-flex-algo].

- Flags: 8-bit flags. Currently all the flags are reserved for future use. They SHOULD be set to zero on transmission and MUST be ignored on receipt.

- Sub-TLVs: optional sub-TLVs to specify the additional attributes of a virtual transport network. Currently no sub-TLV is defined in this document.

3. Advertisement of VTN Topology Attribute

[I-D.dong-lsr-sr-enhanced-vpn] describes the IGP mechanisms to distribute the topology attributes of SR based VTNs. This section describes the BGP-LS mechanism to distribute both the intra-domain and inter-domain topology attributes of SR based VTNs.

3.1. Intra-domain Topology Advertisement

3.1.1. MTR based Topology Advertisement

In section 3.2.1.5 of [RFC7752], Multi-Topology Identifier (MT-ID) TLV is defined, which can contain one or more IS-IS or OSPF Multi-Topology IDs. The MT-ID TLV MAY be present in a Link Descriptor, a Prefix Descriptor, or the BGP-LS Attribute of a Node NLRI.

[I-D.ietf-idr-bgp-ls-segment-routing-ext] defines the BGP-LS extensions to carry the segment routing information using TLVs of BGP-LS Attribute. When MTR is used with SR-MPLS data plane, topology-specific prefix-SIDs and topology-specific Adj-SIDs can be carried in the BGP-LS Attribute associated with the prefix NLRI and link NLRI respectively, the MT-ID TLV is carried in the prefix descriptor and link descriptor to identify the corresponding topology of the SIDs.

[I-D.ietf-idr-bgpls-srv6-ext] defines the BGP-LS extensions to advertise SRv6 segments along with their functions and attributes. When MTR is used with SRv6 data plane, the SRv6 Locator TLV is carried in the BGP-LS Attribute associated with the prefix-NLRI, the MT-ID TLV can be carried in the prefix descriptor to identify the corresponding topology of the SRv6 Locator. The SRv6 End.X SIDs are carried in the BGP-LS Attribute associated with the link NLRI, the MT-ID TLV can be carried in the link descriptor to identify the corresponding topology of the End.X SIDs. The SRv6 SID NLRI is defined to advertise other types of SRv6 SIDs, in which the SRv6 SID Descriptors can include the MT-ID TLV so as to advertise topology-specific SRv6 SIDs.

[RFC7752] also defines the rules of the usage of MT-ID TLV:

"In a Link or Prefix Descriptor, only a single MT-ID TLV containing the MT-ID of the topology where the link or the prefix is reachable is allowed. In case one wants to advertise multiple topologies for a given Link Descriptor or Prefix Descriptor, multiple NLRIs need to be generated where each NLRI contains an unique MT-ID. In the BGP-LS attribute of a Node NLRI, one MT-ID TLV containing the array of MT-IDs of all topologies where the node is reachable is allowed."

Editor's note: the above rules indicates that only one MT-ID is allowed to be carried the Link or Prefix descriptors. When a link or prefix participates in multiple topologies, multiple NLRIs needs to be generated to report all the topologies a link or prefix participates in, together with the topology-specific segment routing information. This would increase the number of BGP Updates and may introduce additional processing burden to both the sending BGP speaker and the receiving network controller. When the number of topologies in a network is not a small number, some optimization may be introduced for the reporting of multi-topology information and the associated segment routing information in BGP-LS. Based on the WG's This will be elaborated in a future version.

3.1.2. Flex-Algo based Topology Advertisement

The Flex-Algo definition [I-D.ietf-lsr-flex-algo] can be used to describe the topological constraints for path computation on a network topology. As specified in [I-D.dong-lsr-sr-enhanced-vpn], the topology of a VTN can be determined by applying Flex-Algo on a default topology.

BGP-LS extensions for Flex-Algo [I-D.ietf-idr-bgp-ls-flex-algo] provide the mechanisms to advertise the Flex-Algo definition information. BGP-LS extensions for SR-MPLS [I-D.ietf-idr-bgp-ls-segment-routing-ext] and SRv6 [I-D.ietf-idr-bgpls-srv6-ext] provide the mechanism to advertise the algorithm-specific segment routing information.

In[I-D.ietf-idr-bgp-ls-segment-routing-ext], algorithm-specific prefix-SIDs can be advertised in BGP-LS attribute associated with Prefix NLRI. In [I-D.ietf-idr-bgpls-srv6-ext], algorithm-specific SRv6 Locators can be advertised in the Prefix NLRI with the SRv6 Locator TLV carried in the associated BGP-LS Attribute, and algorithm-specific End.X SID can be advertised in BGP-LS Attribute associated with the corresponding Link NLRI. Other types of SRv6 SIDs can also be algorithm-specific and are advertised using the SRv6 SID NLRI .

3.2. Inter-Domain Topology Advertisement

[I-D.ietf-idr-bgpls-segment-routing-epe] and [I-D.ietf-idr-bgpls-srv6-ext] defines the BGP-LS extensions for advertisement of BGP topology information between ASes and the BGP Peering Segment Identifiers. Such information could be used by a network controller for the computation and instantiation of inter-AS traffic engineering SR paths.

In some network scenarios, VTNs which span multiple ASes need to be created. The multi-domain VTNs could have different inter-domain connectivity, and may be associated with different set of network resources in each domain and also on the inter-domain links. In order to build the multi-domain VTNs using segment routing, it is necessary to advertise the topology and resource attribute of VTN on the inter-domain links and the associated BGP Peering SIDs.

Depending on the requirement of inter-domain VTNs, different mechanism can be used on the inter-domain connection:

- One EBGP session between two ASes can be established over multiple underlying links. In this case, different underlying links can be used for different inter-domain VTNs which requires link isolation between each other. In another similar case, the EBGP session is established over a single link, while the network resource (e.g. bandwidth) on this link can be partitioned into several pieces, each of which can be considered as a virtual member link. In both cases, different BGP Peer-Adj-SIDs SHOULD be allocated to each underlying physical or virtual member link, and ASBRs SHOULD advertise the VTN identifier associated with each BGP Peer-Adj-SID.

- For inter-domain connection between two ASes, multiple EBGP sessions can be established between different set of peering ASBRs. It is possible that some of these BGP sessions are used for one multi-domain VTN, while some other BGP sessions are used for another multi-domain VTN. In this case, different BGP peer-node-SIDs are allocated to each BGP session, and ASBRs SHOULD advertise the VTN identifier associated with each BGP Peer-node-SIDs.

- At the AS-level topology, different multi-domain VTNs may have different inter-domain connectivity. Different BGP Peer-Set-SIDs can be allocated to represent the groups of BGP peers which can be used for load-balancing in each multi-domain VTN.

In network scenarios where the MT-ID or Flex-Algo is used consistently in multiple ASes covered by a VTN. the approaches to advertise topology-specific BGP peering SIDs are described as below:

- Using MT-based mechanism, the topology-specific BGP peering SIDs can be advertised with the MT-ID associated with the VTN carried in the corresponding link NLRI. This can be supported with the existing mechanisms defined in [RFC7752][I-D.ietf-idr-bgpls-segment-routing-epe] and [I-D.ietf-idr-bgpls-srv6-ext].

o Using Flex-Algo based mechanism, the topology-specific BGP peering SIDs can be advertised together with the admin-group (color) of the corresponding Flex-Algo in the BGP-LS attribute.

In network scenarios where consistent usage of MT-ID or Flex-Algo among multiple ASes can not be expected, then the global-significant VTN-ID can be used to define the AS level topologies. Within each domain, the MT or Flex-Algo based mechanism could still be used for topology advertisement.

3.2.1. VTN ID TLV

A new VTN ID TLV is defined to describe the identifiers of one or more VTNs an intra-domain or inter-domain link belongs to. It can be carried in BGP-LS attribute which is associated with a Link NLRI, or it could be carried as a sub-TLV in the L2 Bundle Member Attribute TLV.

The format of VTN ID TLV is as below:

```
    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           |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |              Flags            |            Reserved           |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                            VTN ID-1                           |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   ~                              ...                              ~
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                            VTN ID-n                           |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Where:

o Type: TBD

o Length: The length of the value field of the sub-TLV. It is variable dependent on the number of VTN IDs included.

o Flags: 16 bit flags. All the bits are reserved, which MUST be set to 0 on transmission and ignored on receipt.

o Reserved: this field is reserved for future use. MUST be set to 0 on transmission and ignored on receipt.

o VTN IDs: One or more 32-bit identifiers to specify the VTNs this link or member link belongs to.

4. Advertisement of VTN Resource Attribute

[I-D.dong-lsr-sr-enhanced-vpn] specifies the mechanism to advertise the resource information associated with each VTN. It is based on the extensions to the advertisement of L2 bundle member links information[RFC8668]. This section defines the corresponding BGP-LS extensions. Two new TLVs are defined to carry the VTN ID and the link attribute flags of either a Layer 3 link or L2 bundle member link. The new VTN ID TLV is defined in section 3.2.1 of this document, and the new Link Attribute Flags TLV is defined in this section. The TE attributes of each Layer 3 link or Layer 2 bundle member link, such as the bandwidth and adj-SIDs, can be advertised using the mechanism as defined in [I-D.ietf-idr-bgp-ls-segment-routin g-ext][I-D.ietf-idr-bgpls-segment-routing-epe] and [I-D.ietf-idr-bgpls-srv6-ext].

4.1. Link Attribute Flags TLV

A new Link attribute Flags TLV is defined to specify the characteristics of a link. It can be carried in BGP-LS attribute which is associated with a Link NLRI, or it could be carried as a sub-TLV in the L2 Bundle Member Attribute TLV. The format of the sub-TLV is as below:

```
    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             |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |             Flags             |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Where:

Type: TBD

Length: 4 octets.

Flags: 16-bit flags. This field is consistent with the Flag field in IS-IS Link Attribute sub-TLV in [RFC5029]. In addition to the flags defined in [RFC5029], A new Flag V is defined in this document. When the V flag is set, it indicates this link is a virtual link.

5. Advertisement of VTN specific Data Plane Identifiers

In network scenarios where each VTN is associated with an independent network topology or Flex-Algo, the topology or Flex-Algo specific SIDs or Locators could be used as the identifier of the VTN in data plane. In network scenarios where multiple VTNs share the same topology or Flex-Algo, additional data plane identifiers would be needed to identify different VTNs.

This section describes the mechanisms to advertise the VTN identifiers with different data plane encapsulations.

5.1. VTN-specific SIDs for SR-MPLS

With SR-MPLS data plane, the VTN identification information is implicitly carried in the SR SIDs of the corresponding VTN. Each node SHOULD allocate VTN-specific Prefix-SIDs for each VTN it participates in. Similarly, VTN-specific Adj-SIDs MAY be allocated for each link which participates in the VTN.

5.1.1. VTN-specific Prefix-SID TLV

A new VTN-specific Prefix-SID TLV is defined to advertise the prefix-SID and its associated VTN. It is derived from VTN specific Prefix-SID sub-TLV of IS-IS [I-D.dong-lsr-sr-enhanced-vpn]. The format of the sub-TLV is as below:

```
 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    |            Flags              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            VTN ID                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      SID/Index/Label(Variable)                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Where:

- Type: TBD

- Length: The length of the value field of the sub-TLV. It is variable dependent on the length of the SID/Index/Label field.

- Flags: 16-bit flags. The high-order 8 bits are the same as in the Adj-SID sub-TLV defined in [RFC8667]. The lower-order 8 bits are reserved for future use, which SHOULD be set to 0 on transmission and MUST be ignored on receipt.

o VTN ID: A 32-bit local identifier to identify the VTN this prefix-SID associates with.

o SID/Index/Label: The same as defined in [RFC8667].

Multiple VTN-specific Prefix-SID TLV could be carried in BGP-LS attribute of the associated Prefix NLRI. The MT-ID in the Prefix descriptors SHOULD be the same as the MT-ID in the definition of these VTNs.

5.1.2. VTN specific Adj-SID TLV

A new VTN-specific Adj-SID TLV is defined to advertise the Adj-SID and its associated VTN. It is derived from VTN specific Adj-SID sub-TLV of IS-IS [I-D.dong-lsr-sr-enhanced-vpn]. The format of the sub-TLV is as below:

```
 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     |            Flags              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                            VTN ID                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      SID/Index/Label(Variable)                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Where:

o Type: TBD

o Length: The length of the value field of the sub-TLV. It is variable dependent on the length of the SID/Index/Label field.

o Flags: 16-bit flags. The high-order 8 bits are the same as in the Adj-SID sub-TLV defined in [RFC8667]. The lower-order 8 bits are reserved for future use, which SHOULD be set to 0 on transmission and MUST be ignored on receipt.

o VTN ID: A 32-bit local identifier to identify the VTN this Adj-SID associates with.

o SID/Index/Label: The same as defined in [RFC8667].

Multiple VTN-specific Adj-SID TLVs MAY be carried in BGP-LS attribute of the associated Link NLRI. The MT-ID in the Link descriptors SHOULD be the same as the MT-ID in the definition of these VTNs.

5.2. VTN-specific Locators for SRv6

With SRv6 data plane, the VTN identification information can be implicitly or explicitly carried in the SRv6 Locator of the corresponding VTN. Network nodes SHOULD allocate VTN-specific Locators for each VTN it participates in. The VTN-specific Locators are used as the covering prefix of VTN-specific SRv6 End SIDs, End.X SIDs and their variants.

Each VTN-specific SRv6 Locator MAY be advertised in a separate Prefix NLRI. If multiple VTNs share the same topology, the topology/ algorithm specific Locator is the covering prefix of a group of VTN-specific Locators. Then the advertisement of VTN-specific locators can be optimized to reduce the amount of information advertised in the control plane. More details about this mechanism will be provided in a future version of this document.

5.3. Dedicated VTN ID in Data Plane

As the number of VTNs increases, some data plane optimization is needed to reduce the amount of SR SIDs and locators allocated for VTNs. Thus a dedicated VTN ID could be encapsulated in the packet as proposed in [I-D.dong-6man-enhanced-vpn-vtn-id].

In that case, the VTN ID encapsulated in data plane can has the same value as the VTN ID in control plane, so that the overhead of advertising the mapping between the VTN IDs in control plane and the corresponding data plane identifiers could be saved.

6. Security Considerations

This document introduces no additional security vulnerabilities to BGP-LS.

The mechanism proposed in this document is subject to the same vulnerabilities as any other protocol that relies on BGP-LS.

7. IANA Considerations

TBD

8. Acknowledgments

The authors would like to thank Shunwan Zhuang and Zhenbin Li for the review and discussion of this document.

9. References

9.1. Normative References

[I-D.dong-spring-sr-for-enhanced-vpn]
Dong, J., Bryant, S., Miyasaka, T., Zhu, Y., Qin, F., and Z. Li, "Segment Routing for Resource Guaranteed Virtual Networks", draft-dong-spring-sr-for-enhanced-vpn-08 (work in progress), June 2020.

[I-D.ietf-idr-bgp-ls-flex-algo]
Talaulikar, K., Psenak, P., Zandi, S., and G. Dawra, "Flexible Algorithm Definition Advertisement with BGP Link-State", draft-ietf-idr-bgp-ls-flex-algo-02 (work in progress), January 2020.

[I-D.ietf-idr-bgp-ls-segment-routing-ext]
Previdi, S., Talaulikar, K., Filsfils, C., Gredler, H., and M. Chen, "BGP Link-State extensions for Segment Routing", draft-ietf-idr-bgp-ls-segment-routing-ext-16 (work in progress), June 2019.

[I-D.ietf-idr-bgpls-segment-routing-epe]
Previdi, S., Talaulikar, K., Filsfils, C., Patel, K., Ray, S., and J. Dong, "BGP-LS extensions for Segment Routing BGP Egress Peer Engineering", draft-ietf-idr-bgpls-segment-routing-epe-19 (work in progress), May 2019.

[I-D.ietf-idr-bgpls-srv6-ext]
Dawra, G., Filsfils, C., Talaulikar, K., Chen, M., daniel.bernier@bell.ca, d., and B. Decraene, "BGP Link State Extensions for SRv6", draft-ietf-idr-bgpls-srv6-ext-02 (work in progress), January 2020.

[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>.

[RFC5029] Vasseur, JP. and S. Previdi, "Definition of an IS-IS Link Attribute Sub-TLV", RFC 5029, DOI 10.17487/RFC5029, September 2007, <https://www.rfc-editor.org/info/rfc5029>.

[RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <https://www.rfc-editor.org/info/rfc7752>.

[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>.

9.2. Informative References

[I-D.dong-6man-enhanced-vpn-vtn-id]
Dong, J. and Z. Li, "Carrying Virtual Transport Network (VTN) Identifier in IPv6 Extensison Header for Enhanced VPN", draft-dong-6man-enhanced-vpn-vtn-id-00 (work in progress), February 2020.

[I-D.dong-lsr-sr-enhanced-vpn]
Dong, J., Hu, Z., Li, Z., and S. Bryant, "IGP Extensions for Segment Routing based Enhanced VPN", draft-dong-lsr-sr-enhanced-vpn-03 (work in progress), March 2020.

[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-07 (work in progress), April 2020.

[I-D.ietf-lsr-isis-srv6-extensions]
Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extension to Support Segment Routing over IPv6 Dataplane", draft-ietf-lsr-isis-srv6-extensions-08 (work in progress), April 2020.

[I-D.ietf-teas-enhanced-vpn]
Dong, J., Bryant, S., Li, Z., Miyasaka, T., and Y. Lee, "A Framework for Enhanced Virtual Private Networks (VPN+) Services", draft-ietf-teas-enhanced-vpn-05 (work in progress), February 2020.

[RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <https://www.rfc-editor.org/info/rfc8667>.

[RFC8668] Ginsberg, L., Ed., Bashandy, A., Filsfils, C., Nanduri, M., and E. Aries, "Advertising Layer 2 Bundle Member Link Attributes in IS-IS", RFC 8668, DOI 10.17487/RFC8668, December 2019, <https://www.rfc-editor.org/info/rfc8668>.

Authors' Addresses

Jie Dong
Huawei Technologies

Email: jie.dong@huawei.com

Zhibo Hu
Huawei Technologies

Email: huzhibo@huawei.com

Zhenbin Li
Huawei Technologies

Email: lizhenbin@huawei.com

Xiongyan Tang
China Unicom

Email: tangxy@chinaunicom.cn

Ran Pang
China Unicom

Email: pangran@chinaunicom.cn