Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 9, 2021

Y. Zhu
China Telecom
R. Chen
S. Peng
ZTE Corporation
F. Qin
China Mobile

September 05, 2020

# IS-IS Extensions to Support Transport Network Slices using Segment Routing
## draft-zch-lsr-isis-network-slicing-06

Abstract

[I-D.nsdt-teas-ns-framework] provides a framework of transport slices.

This draft describes the IS-IS extensions required to support transport slices using Segment Routing.

Table of Contents

## 1. Introduction

[I-D.nsdt-teas-ns-framework] provides a framework for discussing transport slices.

For a Transport Network, transport network slicing requires the underlying network to support partitioning of the network resources to provide the client with dedicated (private) networking, computing, and storage resources drawn from a shared pool. [I-D.peng-teas-network-slicing] introduce a unified TN-slice identifier (termes as AII) to the underlay transport network according to the requirement of network slice. AII is convenient to indicate the topology, computing, storage resources of the dedicated virtual network for both intra-domain and inter-domain network slicing scenarios, and it is flexible to compute SR-BE or SR-TE path according to AII combined with other creteria.

This draft describes the IS-IS extensions required to distribute TN-slice Identifier information in an AS.

## 2. Conventions used in this document

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 RFC2119.

## 3. Router Capabilities for TN-slice Identifier

Although a router can deduce which TN-slices it has participated in according to the AII(administrative instance identifier) configuration of all links, an ISIS instance can explicitly control which TN-slices it wants to enable (or join), to explicitly control which SPT (shortest path tree) for a specific AII to be created. It is possible for a route process not to join any TN-slices (except the default AII 0) in despite of any AII configuration of any links. Especially, it is hard to deduce the participated TN-slice according to the AII configuration of L2 Bundle Member.

This section defines AII Participation sub-TLV which is inserted into the IS-IS Router Capability TLV-242 that is defined in [RFC7981], to explicitly advertise which TN-slice a router wants to take part in.

The TN-slice identifier Participation sub-TLV has the following format:

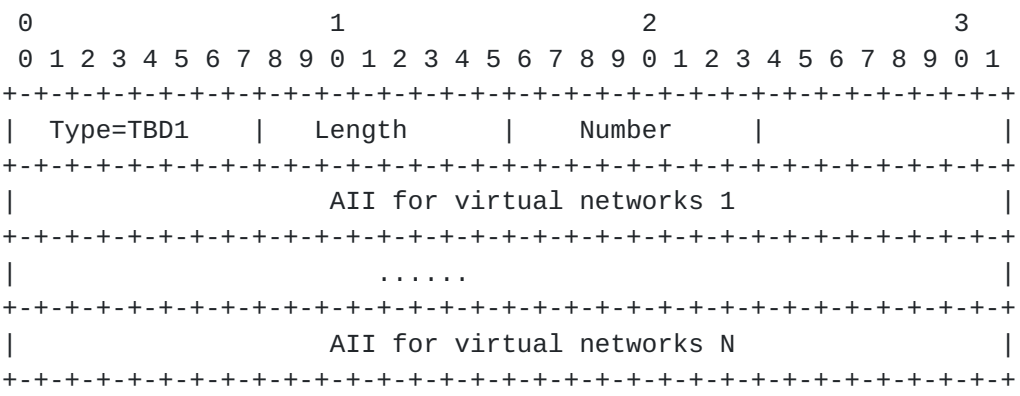


Figure 1: Node Participation to TN-slice

where:

Type: TBD1 (Suggested value to be assigned by IANA)

Length: variable.

Number:Number of virtual networks.

AII for VN: allocate different TN-slice identifier (AII) for different virtual networks. AII is used to distinguish different virtual network resources.

## 4. Advertising TN-slice Identifier as a new TE parameter of a link

[RFC5305] describes extensions to the Intermediate System to Intermediate System (IS-IS) protocol to support Traffic Engineering (TE).

TN-slice Identifier can be used to color links to partition underlay resource. This document defines a new extension of the existing IGP-TE mechanisms [RFC5305] to distribute TN-slice Identifier information in an AS as a new TE parameter of a link.

The TN-slice Identifier list sub-TLV has the following format:

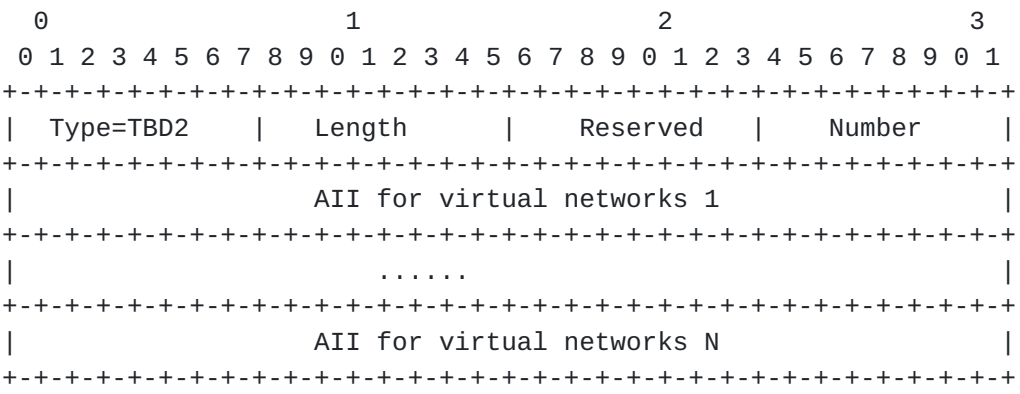


Figure 2: Link Participation to TN-slice

Type: TBD2 (Suggested value to be assigned by IANA)

Length: variable.

Number:Number of virtual networks.

AII for VN: allocate different TN-slice identifier (AII) for different virtual networks. AII is used to distinguish different virtual network resources.

This sub-TLV MAY be present in any of the following TLVs:

TLV-22 (Extended IS reachability) [RFC5305].

TLV-222 (Multitopology IS) [RFC5120].

TLV-23 (IS Neighbor Attribute) [RFC5311].

TLV-223 (Multitopology IS Neighbor Attribute) [RFC5311].

TLV-141 (inter-AS reachability information) [RFC5316].

This sub-TLV SHOULD appear once at most in each TLV. Indicates that a link MAY belong to multiple virtual networks.

Note that AII 0 does not require notification, and all links are always in AII 0 at the same time.

## 4.1. Advertising Dedicated Bandwidth Resouce of a Link for Each VN

[RFC8570] provide ways to distribute network-performance information, and it defines Unidirectional Residual Bandwidth Sub-TLV (URB Sub-TLV), Unidirectional Available Bandwidth Sub-TLV (UAB Sub-TLV), Unidirectional Utilized Bandwidth Sub-TLV (UUB Sub-TLV) for each

directly connected IS-IS neighbors. As described above, a link can belong to multiple VNs, as traditional ways that these VNs can share the total bandwidth resouce of the link with preemption mode based on packet priority, this is what we know as soft slices. However, In some other scenarios, a hard slicing scheme can be used to establish a hardened pipe to meet the slicing business requirements, at this time each VN need dedicated bandwidth resouce reserved from the same link, and in each node the packet rate limit and QoS policy per slice may be used to ensure that the traffic between different slices is isolated and does not affect each other. For this purpose, this section continues to introduce and register new bandwidth related IS-IS TE sub-TLVs in the "Sub-TLVs for TLVs 22, 23, 141, 222, and 223" registry.

### 4.1.1. Unidirectional Residual Bandwidth per TN-slice Identifier Sub-TLV

This sub-TLV (URBPTSI Sub-TLV) advertises the residual bandwidth for a specific slice between two directly connected IS-IS neighbors. The residual bandwidth advertised by this sub-TLV MUST be the residual bandwidth from the system originating the Link State Advertisement (LSA) to its neighbor. This sub-TLV may appear more than once.

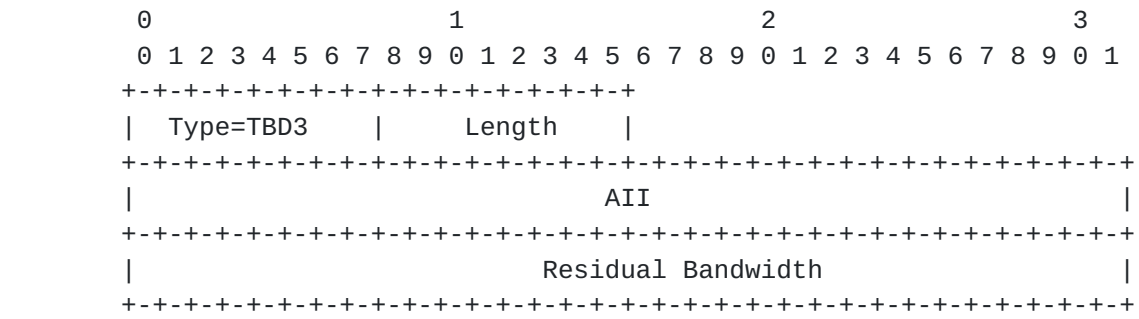


Figure 3: Unidirectional Residual Bandwidth per TN-slice

Type: TBD3 (Suggested value to be assigned by IANA)

Length: 8

AII: Identifies the TN-slice information.

Residual Bandwidth: This field carries the residual bandwidth for specific TN-slice on a link or bundled link in IEEE floating-point format with units of bytes per second.

When AII is 0, the URBPTSI Sub-TLV has the same function with the URB Sub-TLV defined in [RFC8570]. Note that the sum of URBPTSI Sub-TLV for all slices MUST not be larger than the result of Maximum Link Bandwidth [RFC5305] minus RSVP-TE tunnel reservations.

For the system originating the URBPTSI Sub-TLV, the residual bandwidth for each specific slice could be configured explicitly. For example, for a link with 10G bandwidth without any RSVP-TE tunnel reservations, the configuration could be: 4G is allocated for AII 1, 4G is allocated for AII 2, then the rest 2G is allocated for AII 0. Thus, an URBPTSI Sub-TLV for AII 0 or traditional URB Sub-TLV with Residual Bandwidth field set to 2G can be advertised, and two URBPTSI Sub-TLVs each for AII 1 or 2 with Residual Bandwidth field set to 4G can be advertised respectively.

### 4.1.2. Unidirectional Available Bandwidth per TN-slice Identifier Sub-TLV

This sub-TLV (UABPTSI Sub-TLV) advertises the available bandwidth for a specific slice between two directly connected IS-IS neighbors. The available bandwidth advertised by this sub-TLV MUST be the available bandwidth from the system originating this sub-TLV. This sub-TLV may appear more than once. The format of this sub-TLV is shown in the following diagram:

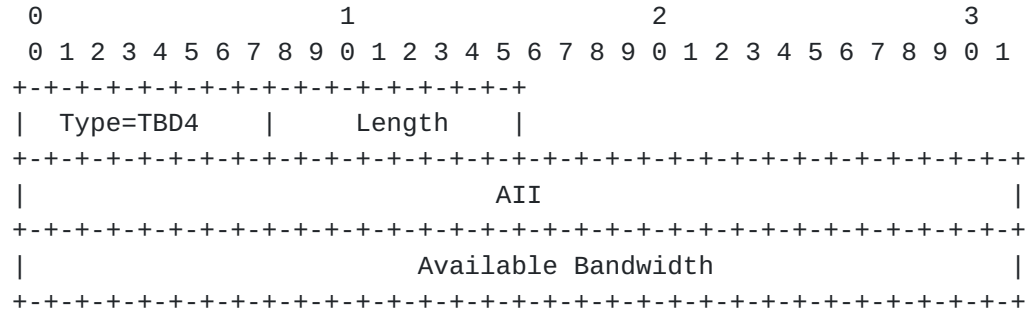


Figure 4: Unidirectional Available Bandwidth per TN-slice

Type: TBD4 (Suggested value to be assigned by IANA)

Length: 8

AII: Identifies the TN-slice information.

Available Bandwidth: This field carries the available bandwidth for specific TN-slice on a link or bundled link in IEEE floating-point format with units of bytes per second. It is defined to be residual

bandwidth per TN-slice minus the measured bandwidth (the UUBPTSI introduced in next section) used for the actual forwarding of packets within that slice.

### 4.1.3. Unidirectional Utilized Bandwidth per TN-slice Identifier Sub-TLV

This sub-TLV (UUBPTSI Sub-TLV) advertises the bandwidth utilization for a specific slice between two directly connected IS-IS neighbors. The bandwidth utilization advertised by this sub-TLV MUST be the bandwidth from the system originating this sub-TLV. This sub-TLV may appear more than once. The format of this sub-TLV is shown in the following diagram:

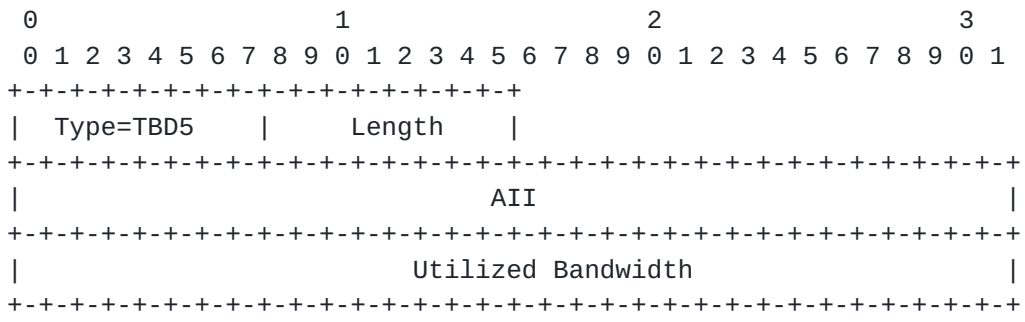


Figure 5: Unidirectional Utilized Bandwidth per TN-slice

Type: TBD5 (Suggested value to be assigned by IANA)

Length: 8

AII: Identifies the TN-slice information.

Utilized Bandwidth: This field carries the bandwidth utilization for specific TN-slice on a link or bundled link in IEEE floating-point format with units of bytes per second. It represents the actual utilization of the link (i.e., as measured by the advertising node) for that slice.

## 5. Advertising TN-slice Identifier for L2 Bundle Member

[RFC8668] defines a sub-TLV of L2 Bundle Attribute Descriptors, and the sub-TLV may define an attribute common to all of the bundle members listed or a sub-TLV may define an attribute unique to each bundle member. For a bundled link, there are two ways to use TN-slice Identifier to partition its resource. The first one is to set the parent interface with one or more expected TN-slice Identifier,

that is we have seen in Section 4. The second one is to set each member interface with specific TN-slice Identifier, for example, member 1 set to AII 0, member 2 set to AII 1, member 3 set to AII 2, etc. For simple deployment and according to actual needs, each member is set to a single TN-slice Identifier. For the second case, this section defines a new sub-TLV: L2 Bundle Member TN-slice Identifier sub-TLV, to advertise TN-slice Identifier for each L2 Bundle Member associated with a parent L3 adjacency which is Point-to-Point. The following format is defined for this sub-TLV:

Type: TBD6.

Length: variable

L2 Bundle Member TN-slice Identifier. There MUST be one TN-slice Identifier(AII) for each of the L2 Bundle Members advertised under the preceding L2 Bundle Member Attribute Descriptor.

This sub-TLV MAY be present in the following TLVs:

TLV-25 (L2 Bundle Member Attributes) [RFC8668].

This sub-TLV SHOULD appear once at most in the TLV. Indicates that the TN-slice Identifier for L2 Bundle Member.

## 5.1. Advertising Bandwidth Information of member Link

[RFC8668] allows notification of bandwidth information of member links using the sub-TLVs defined in [RFC8570]. Because each member link belongs to single TN-slice, it is not necessary to advertise the bandwidth information of member links for specific slice using the sub-TLVs defined in Section 4.1.

## 6. Advertising prefix-SID per TN-slice Identifier

[RFC8667] defines a new IS-IS sub-TLV: the Prefix Segment Identifier sub-TLV (Prefix-SID sub-TLV). The Prefix-SID sub-TLV carries the Segment Routing IGP-Prefix-SID as defined in [RFC8402], and is associated to a prefix advertised by a node.

To distinguish forwarding behavior of different virtual networks, Prefix-SID need to be allocated per TN-slice Identifier and advertised in the IGP domain. This document defines a new extension of the existing Prefix-SID sub-TLV.

The Prefix-SID for TN-slice Identifier sub-TLV has the following format:

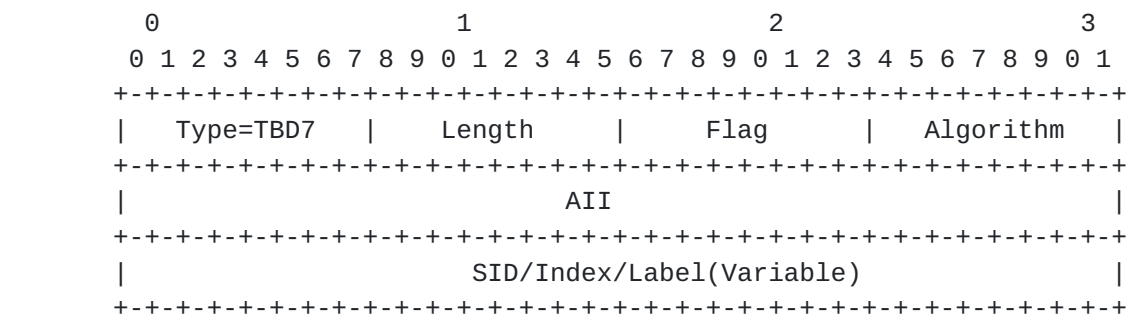


Figure 6: Prefix-SID per TN-slice

where:

Type: TBD7 (Suggested value to be assigned by IANA)

Length: Variable. Depending on the size of the SID.

The "Flags" and "SID/Index/Label" fields are the same as the Prefix-SID sub-TLV [RFC8667] .

Algorithm: Accoridng to section "3.2. SR-Algorithm Sub-TLV" of [RFC8667] , two values can be set in this field.

- 0: Shortest Path First (SPF) algorithm based on link metric.
- 1: Strict Shortest Path First (SPF) algorithm based on link metric.

Note that[I-D.ietf-lsr-flex-algo].also allows user to define other algorithm values, i.e., FA-id within [128, 255], for the purpose of constraint based path computation. However, an FA-id algorithm value MUST not be set in this field, the reason is that FA-id has not semantic local within AII.

AII: Identifies the TN-slice (AII) information corresponding to the Prefix-SID.

This sub-TLV MAY be present in any of the following TLVs:

TLV-135 (Extended IPv4 reachability) defined in [RFC5305].

TLV-235 (Multitopology IPv4 Reachability) defined in [RFC5120].

TLV-236 (IPv6 IP Reachability) defined in [RFC5308].

TLV-237 (Multitopology IPv6 IP Reachability) defined in [RFC5120].

This sub-TLV MAY appear multiple times in each TLV.

## 7. Advertising Adjacency-SID per TN-slice Identifier

[RFC8667] defines the IS-IS sub-TLV: the Adjacency Segment Identifier sub-TLV (Adj-SID sub-TLV).The Adj-SID sub-TLV is an optional sub-TLV carrying the Segment Routing IGP-Adjacency-SID as defined in [RFC8402].

To distinguish forwarding behavior of different virtual networks, Adjacency-SID need to be allocated per TN-slice Identifier and advertised in the IGP domain. This document defines a new extension of the existing Adjacency-SID sub-TLV.

The Adjacency-SID for TN-slice Identifier sub-TLV has the following format:

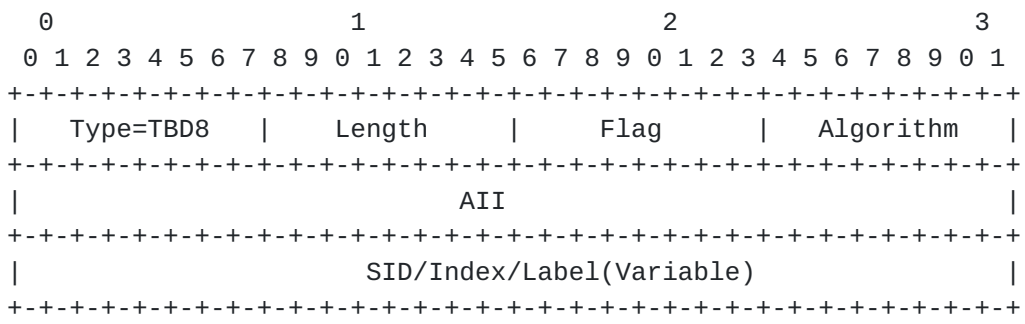


Figure 7: Adjacency-SID per TN-slice

where:

Type: TBD8 (Suggested value to be assigned by IANA)

Length: Variable. Depending on the size of the SID.

The "Flags" and "SID/Index/Label" fields are the same as the Adjacency-SID sub-TLV [RFC8667].

AII: Identifies the TN-slice (AII) information corresponding to the Adjacency-SID.

This sub-TLV MAY be present in any of the following TLVs:

TLV-22 (Extended IS reachability) [RFC5305].

TLV-222 (Multitopology IS)[RFC5120].

TLV-23 (IS Neighbor Attribute)[RFC5311].

TLV-223 (Multitopology IS Neighbor Attribute)[RFC5311].

TLV-141 (inter-AS reachability information)[RFC5316].

Multiple Adj-SID sub-TLVs MAY be associated with a single IS-neighbor. This sub-TLV MAY appear multiple times in each TLV.

## 8. Advertising Adjacency-SID per TN-slice Identifier in LANs

In LAN subnetworks,[RFC8667] defines the LAN-Adj-SID sub-TLV for a router to advertise the Adj-SID of each of its neighbors.

To distinguish forwarding behavior of different virtual networks, Adjacency-SID need to be allocated per TN-slice Identifier and advertised in the IGP domain. This document defines a new extension of the existing Adjacency-SID sub-TLV.

The LAN-Adj-SID for TN-slice Identifier sub-TLV has the following format:

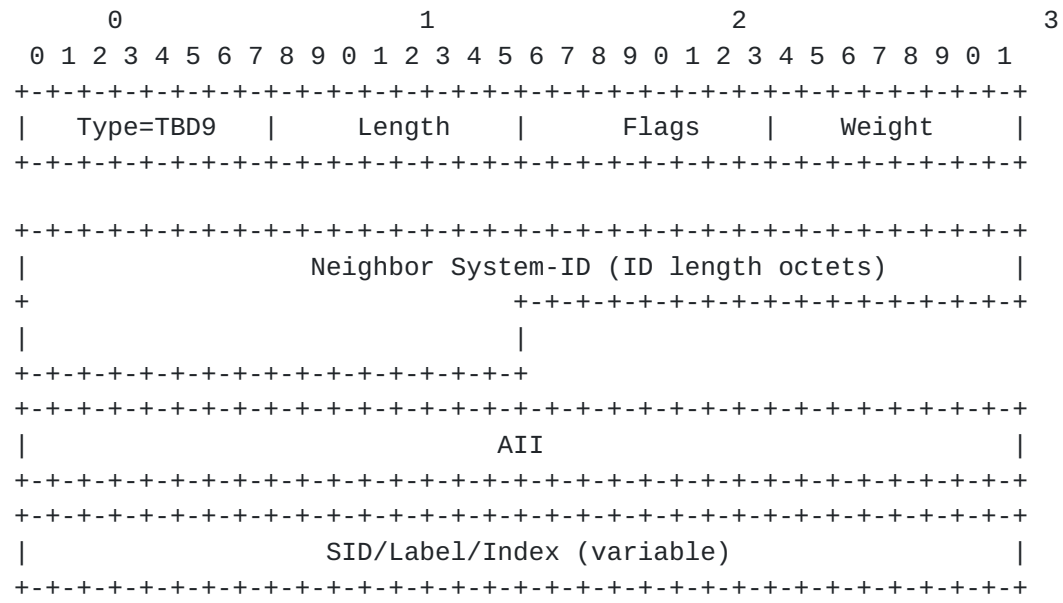


Figure 8: Adjacency-SID per TN-slice in LANs

where:

Type: TBD9 (Suggested value to be assigned by IANA)

Length: Variable. Depending on the size of the SID.

The "Flags" and "SID/Index/Label" fields are the same as the Adjacency-SID sub-TLV [RFC8667].

AII: Identifies the TN-slice (AII) information corresponding to the Adjacency-SID.

This sub-TLV MAY be present in any of the following TLVs:

TLV-22 (Extended IS reachability) [RFC5305].

TLV-222 (Multitopology IS)[RFC5120].

TLV-23 (IS Neighbor Attribute)[RFC5311].

TLV-223 (Multitopology IS Neighbor Attribute)[RFC5311].

Multiple Adj-SID sub-TLVs MAY be associated with a single IS-neighbor. This sub-TLV MAY appear multiple times in each TLV.

## 9. IANA Considerations

This document requests allocation for the following Sub-TLVs.

### 9.1. Router Capabilities for TN-slice Identifier

This document requests IANA to assign a new code point in the "sub-TLV for TLV 242" registry.

Type: TBD1

### 9.2. TN-slice Identifier list sub-TLV

This TLV shares sub-TLV space with existing "Sub-TLVs for TLVs 22, 222, 23, 223 and 141 registry".

Type:TBD2(to be assigned by IANA).

### 9.3. Unidirectional Residual Bandwidth per TN-slice Identifier Sub-TLV

This TLV shares sub-TLV space with existing "Sub-TLVs for TLVs 22, 222, 23, 223 and 141 registry".

Type:TBD3(to be assigned by IANA).

### 9.4. Unidirectional Available Bandwidth per TN-slice Identifier Sub-TLV

This TLV shares sub-TLV space with existing "Sub-TLVs for TLVs 22, 222, 23, 223 and 141 registry".

Type:TBD4(to be assigned by IANA).

### 9.5. Unidirectional Utilized Bandwidth per TN-slice Identifier Sub-TLV

This TLV shares sub-TLV space with existing "Sub-TLVs for TLVs 22, 222, 23, 223 and 141 registry".

Type:TBD5(to be assigned by IANA).

### 9.6. L2 Bundle Member TN-slice Identifier sub-TLV

This TLV shares sub-TLV space with existing "Sub-TLVs for TLVs 22, 222, 23, 25, 223 and 141 registry.

Type:TBD6(to be assigned by IANA).

This sub-TLV is allowed in the following TLVs:

| 22 | 23 | 25 | 141 | 222 | 223 |
|---|---|---|---|---|---|
| n | n | y | n | n | n |

### 9.7. Prefix-SID for TN-slice Identifier sub-TLV

This TLV shares sub-TLV space with existing "Sub-TLVs for TLVs 135,235,226 and 237 registry".

Type:TBD7(to be assigned by IANA).

### 9.8. Adjacency-SID for TN-slice Identifier sub-TLV

This TLV shares sub-TLV space with existing "Sub-TLVs for TLVs 22, 222, 23, 223 and 141 registry".

Type:TBD8(to be assigned by IANA).

### 9.9. LAN-Adj-SID for TN-slice Identifier sub-TLV

This TLV shares sub-TLV space with existing "Sub-TLVs for TLVs 22, 222, 23, and 223 registry".

Type:TBD9(to be assigned by IANA).

## 10. Security Considerations

TBD.

## 11. Acknowledgements

TBD.

## 12. References

### 12.1. Normative references

[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-10 (work in progress), August 2020.

[I-D.nsdt-teas-ns-framework]
Gray, E. and J. Drake, "Framework for Transport Network Slices", draft-nsdt-teas-ns-framework-04 (work in progress), July 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>.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <https://www.rfc-editor.org/info/rfc5120>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <https://www.rfc-editor.org/info/rfc5305>.

[RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <https://www.rfc-editor.org/info/rfc5308>.

[RFC5311] McPherson, D., Ed., Ginsberg, L., Previdi, S., and M. Shand, "Simplified Extension of Link State PDU (LSP) Space for IS-IS", RFC 5311, DOI 10.17487/RFC5311, February 2009, <https://www.rfc-editor.org/info/rfc5311>.

[RFC5316] Chen, M., Zhang, R., and X. Duan, "ISIS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 5316, DOI 10.17487/RFC5316, December 2008, <https://www.rfc-editor.org/info/rfc5316>.

[RFC7981] Ginsberg, L., Previdi, S., and M. Chen, "IS-IS Extensions for Advertising Router Information", RFC 7981, DOI 10.17487/RFC7981, October 2016, <https://www.rfc-editor.org/info/rfc7981>.

[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>.

[RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 8570, DOI 10.17487/RFC8570, March 2019, <https://www.rfc-editor.org/info/rfc8570>.

[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>.

## 12.2. Informative references

[I-D.peng-teas-network-slicing] Peng, S., Chen, R., Mirsky, G., and F. Qin, "Packet Network Slicing using Segment Routing", draft-peng-teas-network-slicing-03 (work in progress), February 2020.

Authors' Addresses

Yongqing Zhu
China Telecom

Email: zhuyq.gd@chinatelecom.cn

Ran Chen
ZTE Corporation

Email: chen.ran@zte.com.cn

Shaofu Peng
ZTE Corporation

Email: peng.shaofu@zte.com.cn

Fengwei Qin
China Mobile

Email: qinfengwei@chinamobile.com