IDR
Internet-Draft
Intended status: Standards Track
Expires: April 16, 2021

R. chen
sh. peng
ZTE Corporation

October 13, 2020

BGP-LS Extensions for Transport Slice over IPv6 Dataplane
draft-chen-idr-bgp-ls-transport-slice-srv6-01

Abstract

[I-D.peng-teas-network-slicing]defines a unified TN-slice identifier, AII(administrative instance identifier), to indicate the topology, computing, storage resources of the dedicated virtual network for both intra-domain and inter-domain network slicing scenarios. This draft defines extensions to BGP-LS protocol in order to advertise the information of the transport slice over IPv6 dataplane..

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 April 16, 2021.

Table of Contents

1. Introduction

For a packet network, 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]defines a unified TN-slice identifier, AII(administrative instance identifier), to indicate the topology, computing, storage resources of the dedicated virtual network for both intra-domain and inter-domain network slicing scenarios, and how to compute SR-BE or SR-TE path according to TN-slice Identifier combined with other criteria.

[I-D.peng-lsr-isis-network-slicing-srv6]describes the ISIS extensions required to support Packet Network Slicing over IPv6 dataplane.

For SRv6 case, IPv6 address resource is directly used to represent SID, so that different IPv6 block could be allocated to different TN-slice. There are two possible ways to advertise TN-slice specfic IPv6 block:

- Traditional prefix reachability, for default AII (0) specific IPv6 block.

- New SRv6 Locator per TN-slice advertisement, for nonzero TN-slice specific IPv6 block.

In order to satisfy the need for applications that require topological visibility across one area or Autonomous System (AS). This document specifies extensions to the BGP Link-state address-family in order to advertise TN-slice over IPv6 dataplane. An external component (e.g., a controller) then can collect TN-slice information in the "northbound" direction.

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. BGP-LS Extensions for transport slice over over IPv6 dataplane

The following transport slice information TLV is defined: BGP-LS[[RFC7752]]defines the link-state NLRI that can be a Node NLRI, a Link NLRI or a Prefix NLRI. BGP-LS[[RFC7752]]defines the TLVs that map link-state information to BGP link-state NLRI within the BGP-LS Attribute. This document adds additional BGP-LS Attribute TLVs in order to encode TN-slice information. It does not introduce any changes to the encoding of the BGP-LS NLRIs.

4. Node Attributes

4.1. SRv6 Node MSD Types

The SRv6 Node MSD Types as defined in [I-D.ietf-idr-bgpls-srv6-ext].

4.2. Router Capabilities for TN-slice Identifier

This BGP-LS Attribute TLV is used to announce which TN-slice a router wants to take part in. This TLV maps to the TN-slice identifier Participation sub-TLV [I-D.zch-lsr-isis-network-slicing] of the IS-IS protocol.

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

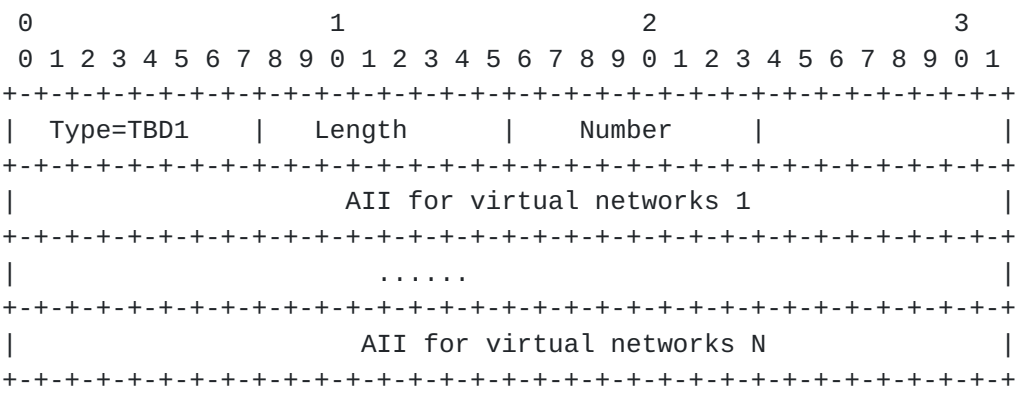

Figure 1

where:

Type: TBA1 (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.

5. Link Attributes

The following Link Attribute TLVs are defined:

Type	Description
TBA1	SRv6 End.X SID per TN-slice TLV
TBA2	SRv6 LAN End.X SID per TN-slice sub-TLV

Table 1:The new Link Attribute TLVs

These TLVs should only be added to the BGP-LS Attribute associated with the Link NLRI.

5.1. SRv6 End.X SID per TN-slice TLV

This sub-TLV is used to advertise a TN-slice specific SRv6 SID associated with a point to point adjacency. Multiple SRv6 End.X SID per TN-slice sub-TLVs MAY be associated with the same adjacency.

The SRv6 End.X SID per TN-slice sub-TLV has the following format:

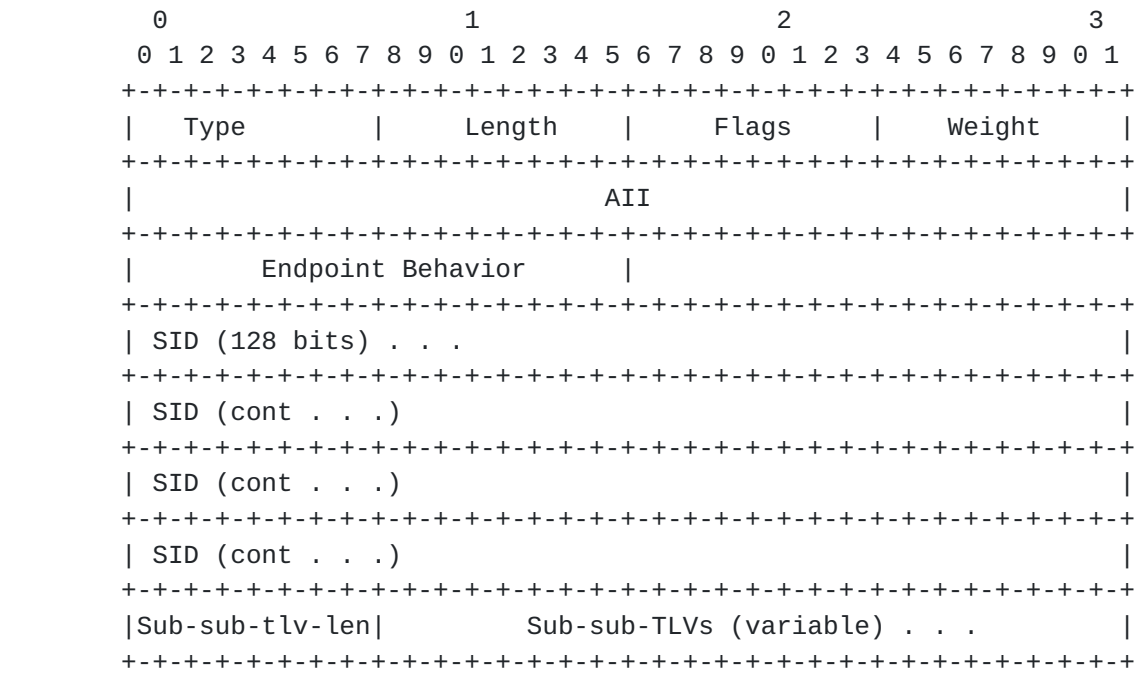

Figure 2

Type: TBA2 (Suggested value to be assigned by IANA)

Length: variable.

Flags: 1 octet.

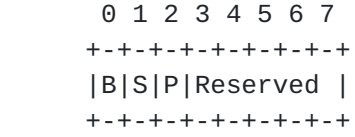

B-Flag:Backup flag. If set, the End.X SID is eligible for protection (e.g., using IPFRR) as described in [RFC8355].

S-Flag:Set flag. When set, the S-Flag indicates that the End.X SID refers to a set of adjacencies (and therefore MAY be assigned to other adjacencies as well).

P-Flag:Persistent flag. When set, the P-Flag indicates that the End.X SID is persistently allocated, i.e., the End.X SID value remains consistent across router restart and/or interface flap.

Other bits: MUST be zero when originated and ignored when received.

Weight: 1 octet. The value represents the weight of the End.X SID for the purpose of load balancing. The use of the weight is defined in [RFC8402].

AII: 4 octet. Administrative Instance Identifier, As defined in [I-D.peng-lsr-isis-network-slicing-srv6], represented as TN-slice Identifier.

Endpoint Behavior: 2 octets. As defined in [I-D.ietf-spring-srv6-network-programming].Legal behavior values for this sub-TLV are defined in Section 8.

SID: 16 octets. This field encodes the advertised SRv6 SID.

Sub-sub-TLV-length: 1 octet. Number of octets used by sub-sub-TLVs

Note that multiple TLVs for the same neighbor may be required in order to advertise all of the SRv6 End.X SIDs associated with that neighbor.

5.2. SRv6 LAN End.X SID per TN-slice TLV

This sub-TLV is used to advertise a TN-slice specific SRv6 SID associated with a LAN adjacency. Since the parent TLV is advertising an adjacency to the Designated Intermediate System(DIS) for the LAN, it is necessary to include the System ID of the physical neighbor on the LAN with which the SRv6 SID is associated. Given that a large number of neighbors may exist on a given LAN a large number of SRv6 LAN END.X SID per TN-slice sub-TLVs may be associated with the same LAN. Note that multiple TLVs for the same DIS neighbor may be required in order to advertise all of the TN-slice specific SRv6 End.X SIDs associated with that neighbor.

The SRv6 LAN End.X SID per TN-slice sub-TLV has the following format:

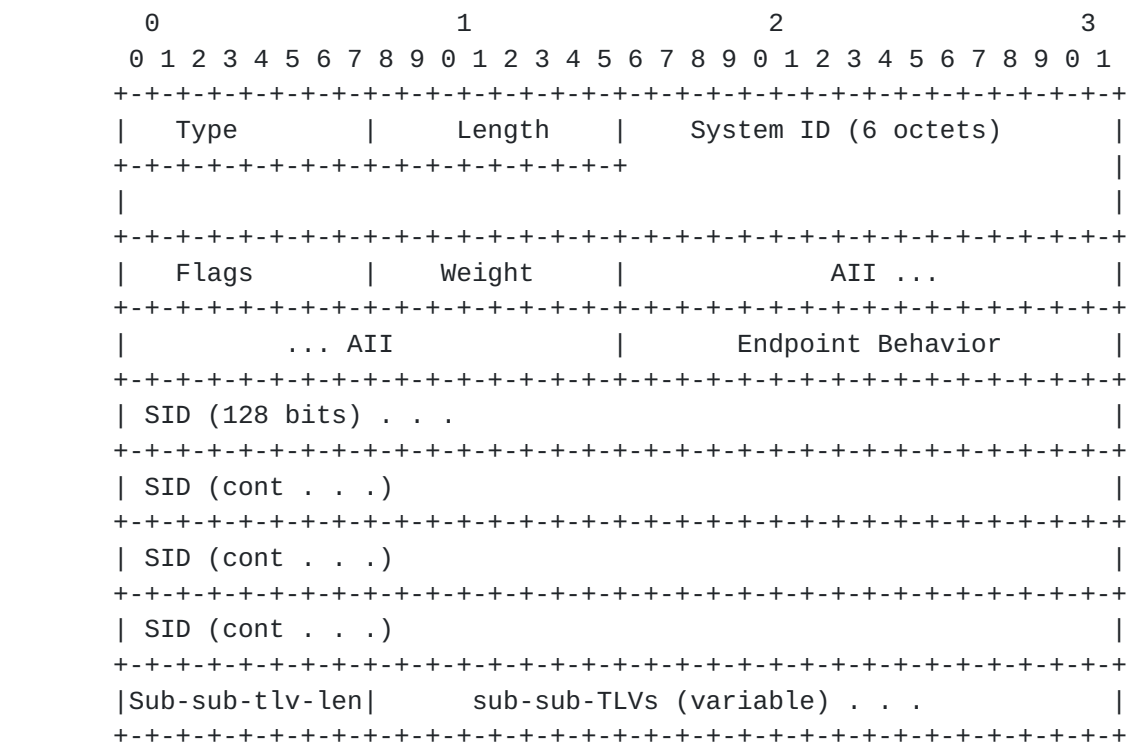

Figure 3

Type: TBA3.

Length: variable

System-ID: 6 octets of IS-IS System-ID of length "ID Length" as defined in [ISO10589].

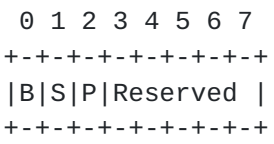

where B,S, and P flags are as described in Section 7.1. Other bits: MUST be zero when originated and ignored when received.

Weight: 1 octet. The value represents the weight of the End.X SID for the purpose of load balancing. The use of the weight is defined in [RFC8402].

AII: 4 octet. Administrative Instance Identifier, As defined in [I-D.peng-lsr-isis-network-slicing-srv6], represented as TN-slice Identifier.

Endpoint Behavior: 2 octets. As defined in [I-D.ietf-spring-srv6-network-programming].Legal behavior values for this sub-TLV are defined in Section 8.

SID: 16 octets. This field encodes the advertised SRv6 SID.

Sub-sub-TLV-length: 1 octet. Number of octets used by sub-sub-TLVs

6. Prefix Attributes

SRv6 attributes with an IPv6 prefix are advertised using the new BGP-LS Attribute TLVs defined in this section and associated with the BGP-LS Prefix NLRI.

6.1. SRv6 Locator per TN-slice TLV

The SRv6 Locator per TN-slice TLV is introduced to advertise SRv6 Locators and End SIDs associated with each locator for specific TN-slice.

The SRv6 Locator per TN-slice TLV has the following format:

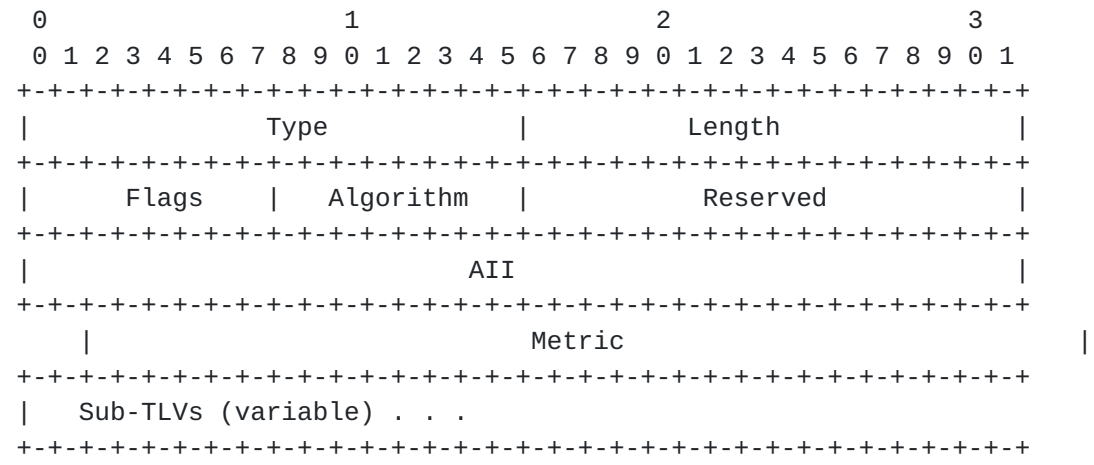

Figure 4

Type:TBA4(Suggested value to be assigned by IANA).

Length:2 octet field with the total length of the value portion of the TLV.

Flags: 1 octet of flags with the following definition:

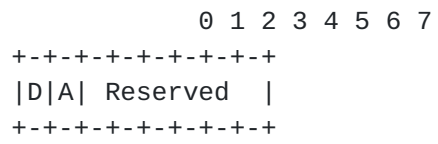

D-Flag: Indicates that the locator has been leaked into the IGP domain when set. IS-IS operations for this are discussed in[I-D.ietf-isis-segment-routing-extensions].

Reserved bits: Reserved for future use and MUST be zero when originated and ignored when received.

Algorithm: 1 octet field. Algorithm associated with the SID. Algorithm values are defined in the IGP Algorithm Type registry.

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

Metric: 4 octet field. The value of the metric for the Locator.

Sub-TLVs : currently none defined. Used to advertise sub-TLVs that provide additional attributes for the given SRv6 Locator.

7. SRv6 SID NLRI

The SRv6 SID NLRI defined in [I-D.ietf-idr-bgpls-srv6-ext] can be reused in this document to advertise TN-slice specific SRv6 Segment Identifiers (SID).

7.1. SRv6 SID per specific TN-slice

A SRv6 SID is a 128 bit value [I-D.ietf-spring-srv6-network-programming] and AII Identifies the TN-slice are encoded using the SRv6 SID for specific TN-slice TLV.

The TLV has the following format:

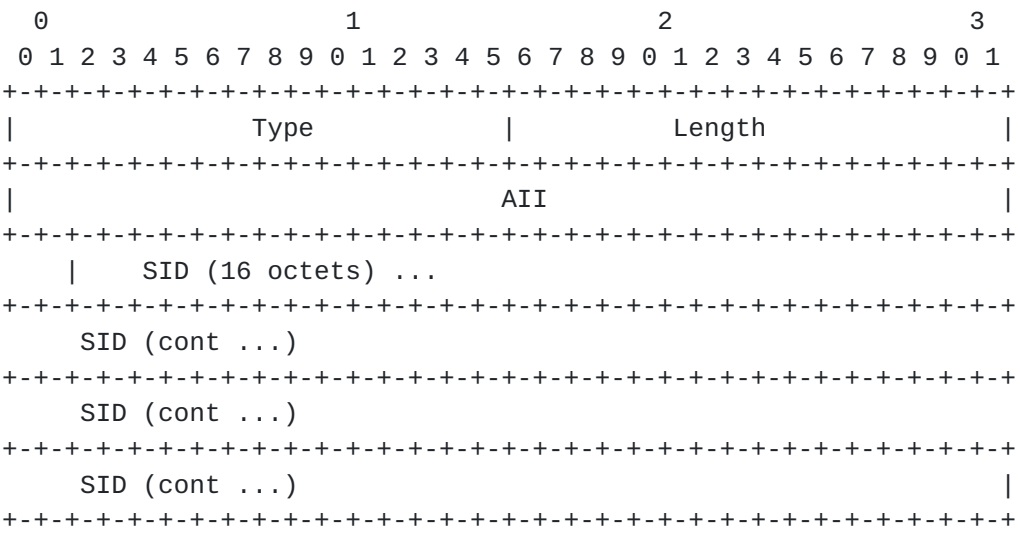

Figure 5

Where:

Type:TBA4(Suggested value to be assigned by IANA).

Length: 2 octet field.

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

SID: 16 octet field. This field encodes the advertised SRv6 SID as 128 bit value.

7.2. SID Attributes

The SRv6 Endpoint Behavior TLV, SRv6 BGP Peer Node SID TLV, and SRv6 SID Structure TLV which are defined in [I-D.ietf-idr-bgpls-srv6-ext]can be reused in this document.

8. IANA Considerations

This document requests assigning code-points from the IANA "Border Gateway Protocol - Link State (BGP-LS) Parameters" registry as described in the sub-sections below.

8.1. BGP-LS TLVs

The following TLV codepoints are assigned by IANA via the early allocation process from within the sub-registry called "BGP-LS Node Descriptor, Link Descriptor, and Prefix Descriptor":

TLV Code Point	Description	Value defined in
TBA1	TN-slice identifier Participation TLV	this document
TBA2	SRv6 End.X SID per TN-slice TLV	this document
TBA3	SRv6 LAN End.X SID per TN-slice sub-TLV	this document
TBA4	SRv6 Locator per TN-slice TLV	this document
TBA5	SRv6 SID per specific TN-slice	this document

9. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the 'Security Considerations' section of [RFC4271] for a discussion of BGP security. Also refer to [RFC4272] and [RFC6952] for analysis of security issues for BGP.

10. Acknowledgements

TBD

11. Normative References

[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-03 (work in progress), July 2020.

[I-D.ietf-isis-segment-routing-extensions]
Previdi, S., Ginsberg, L., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", draft-ietf-isis-segment-routing-extensions-25 (work in progress), May 2019.

[I-D.ietf-spring-srv6-network-programming]
Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 Network Programming", draft-ietf-spring-srv6-network-programming-24 (work in progress), October 2020.

[I-D.peng-lsr-isis-network-slicing-srv6]
Peng, S., Chen, R., and G. Mirsky, "ISIS Extension to Support Network Slicing over IPv6 Dataplane", draft-peng-lsr-isis-network-slicing-srv6-00 (work in progress), January 2020.

[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.

[I-D.zch-lsr-isis-network-slicing]
Zhu, Y., Chen, R., Peng, S., and F. Qin, "IS-IS Extensions to Support Transport Network Slices using Segment Routing", draft-zch-lsr-isis-network-slicing-06 (work in progress), September 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>.

[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>.

[RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <https://www.rfc-editor.org/info/rfc4272>.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <https://www.rfc-editor.org/info/rfc4915>.

[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>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, DOI 10.17487/RFC6952, May 2013, <https://www.rfc-editor.org/info/rfc6952>.

[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>.

[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>.

[RFC8355] Filsfils, C., Ed., Previdi, S., Ed., Decraene, B., and R. Shakir, "Resiliency Use Cases in Source Packet Routing in Networking (SPRING) Networks", RFC 8355, DOI 10.17487/RFC8355, March 2018, <https://www.rfc-editor.org/info/rfc8355>.

[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>.

Authors' Addresses

Ran Chen
ZTE Corporation
No.50 Software Avenue, Yuhuatai District
Nanjing
China

Email: chen.ran@zte.com.cn

Shaofu Peng
ZTE Corporation
No.68 Zijinghua Road, Yuhuatai District
Nanjing
China

Email: peng.shaofu@zte.com.cn