

IDR
Internet-Draft
Intended status: Standards Track
Expires: 23 March 2024

R. Chen
Sh. Peng
ZTE Corporation
P. Psenak
Cisco Systems
L. Yao
ZTE Corporation
19 September 2023

Algorithm Related Adjacency SID Advertisement
draft-chen-idr-bgp-ls-algo-related-adjacency-sid-03

Abstract

This draft describes a mechanism to distribute SR-MPLS algorithm Related Adjacency SID from controllers to the headend routers using BGP-LS.

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 25 March 2024.

Copyright Notice

Copyright (c) 2023 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 Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1.](#) Introduction [2](#)
- [2.](#) Requirements Language [3](#)
- [3.](#) BGP-LS Extensions for SR-MPLS Adjacency with optional [3](#)
 - [3.1.](#) BGP-LS Extensions for link Attribute [3](#)
 - [3.1.1.](#) Adjacency SID per Algorithm TLV [3](#)
 - [3.1.2.](#) LAN Adjacency SID per Algorithm TLV [4](#)
 - [3.2.](#) BGP-LS Extensions for Segment Descriptors [6](#)
 - 3.2.1. Type TBD3: IPv4 Address and Local Interface ID with optional Algorithm [6](#)
 - 3.2.2. Type TBD4: IPv4 Addresses for link endpoints as Local, Remote pair with optional Algorithm [7](#)
 - 3.2.3. Type TBD5: IPv6 Prefix and Interface ID for link endpoints as Local, Remote pair, with optional Algorithm for SR-MPLS [8](#)
 - 3.2.4. Type TBD6: IPv6 Addresses for link endpoints as Local, Remote pair, with optional Algorithm for SR-MPLS [9](#)
- [4.](#) Acknowledgements [10](#)
- [5.](#) IANA Considerations [10](#)
 - [5.1.](#) BGP-LS Link Attribute TLVs registry [10](#)
 - [5.2.](#) BGP-LS SR Segment Descriptors [10](#)
- [6.](#) Security Considerations [10](#)
- [7.](#) Normative References [10](#)
- Authors' Addresses [11](#)

1. Introduction

[I-D.ietf-lsr-algorithm-related-adjacency-sid] defines the algorithm identifier can be included as part of an Adjacency-SID advertisement for SR-MPLS.

[I-D.peng-idr-segment-routing-te-policy-attr] proposes extensions of BGP and defines new Segment Types to provide algorithm information for SR-MPLS Adjacency-SIDs when delivering SR Policy via BGP.

[RFC9085] defines extensions to the Border Gateway Protocol-Link State (BGP-LS) address family in order to carry SR information via BGP, and an algorithm identifier is included as part of the Prefix-SID TLV advertisement.

This draft describes a mechanism to distribute SR-MPLS algorithm Related Adjacency SID from controllers to the headend routers using BGP-LS.

2. 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].

3. BGP-LS Extensions for SR-MPLS Adjacency with optional

3.1. BGP-LS Extensions for link Attribute

[I-D.ietf-lsr-algorithm-related-adjacency-sid] defines the algorithm identifier can be included as part of an Adjacency-SID advertisement for SR-MPLS. This section defines the encoding of SR-MPLS algorithm Related Adjacency SID describes in [I-D.ietf-lsr-algorithm-related-adjacency-sid].

The following link attribute TLVs are defined:

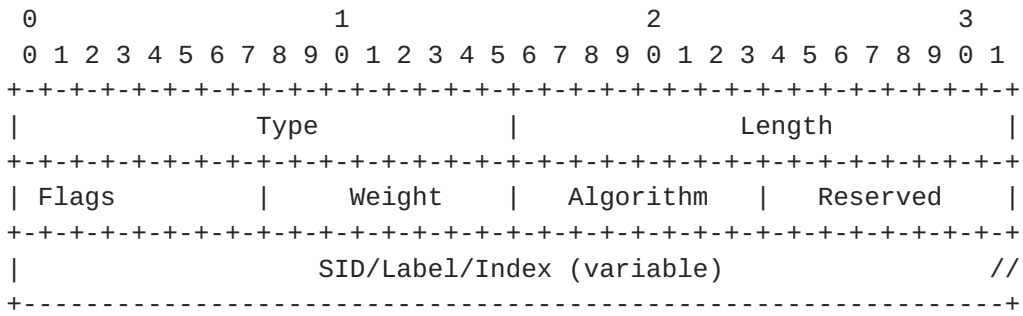
Type	Description
TBD1	Adjacency SID per Algorithm
TBD2	LAN Adjacency SID per Algorithm

Table 1: New link attribute TLVs

3.1.1. Adjacency SID per Algorithm TLV

The Adjacency SID per Algorithm TLV is used in order to advertise information related to an algorithm Related Adjacency SID. This information is derived from the Adjacency Segment Identifier (Adj-SID) per Algorithm Sub-TLV of ISIS/OSPFv2/OSPFv3(Section 4 of [I-D.ietf-lsr-algorithm-related-adjacency-sid]).

The Adjacency SID per Algorithm TLV has the following format:



where:

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

Length: variable.

Flags: 1-octet value that should be set as:

- * IS-IS Adj-SID per Algorithm flags as defined in Section 4.1.1 of [[I-D.ietf-lsr-algorithm-related-adjacency-sid](#)]
- * OSPFv2 Adj-SID per Algorithm Sub-TLV as defined in Section 4.2.1 of [[I-D.ietf-lsr-algorithm-related-adjacency-sid](#)]
- * OSPFv3 Adj-SID per Algorithm Sub-TLV as defined in Section 4.3.1 of [[I-D.ietf-lsr-algorithm-related-adjacency-sid](#)]

Weight: 1 octet carrying the weight used for load-balancing purposes. The use of weight is described in [Section 3.4 of \[RFC8402\]](#).

Algorithm: Refer to ISIS/OSPFv2/OSPFv3 Adj-SID per Algorithm([\[I-D.ietf-lsr-algorithm-related-adjacency-sid\]](#)).

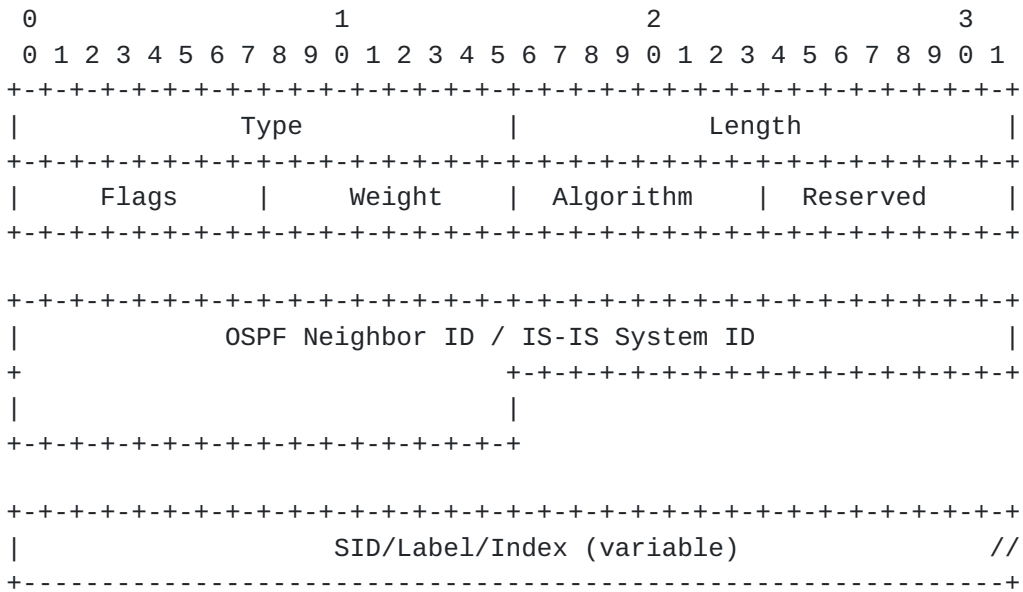
Reserved: 1-octet that MUST be set to 0 and ignored on receipt.

SID/Label/Index:Refer to ISIS/OSPFv2/OSPFv3 Adj-SID per Algorithm([\[I-D.ietf-lsr-algorithm-related-adjacency-sid\]](#)).

3.1.2. LAN Adjacency SID per Algorithm TLV

This information is derived from LAN-Adj-SID per Algorithm Sub-TLV of ISIS/OSPFv2/OSPFv3(Section 4 of [[I-D.ietf-lsr-algorithm-related-adjacency-sid](#)]).

The LAN Adjacency SID per Algorithm TLV has the following format:



where:

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

Length: variable.

Flags: 1-octet value that should be set as:

- * IS-IS LAN-Adj-SID per Algorithm flags as defined in Section 4.1.2 of [\[I-D.ietf-lsr-algorithm-related-adjacency-sid\]](#)
- * OSPFv2 LAN-Adj-SID per Algorithm Sub-TLV as defined in Section 4.2.2 of [\[I-D.ietf-lsr-algorithm-related-adjacency-sid\]](#)
- * OSPFv3 LAN-Adj-SID per Algorithm Sub-TLV as defined in Section 4.3.2 of [\[I-D.ietf-lsr-algorithm-related-adjacency-sid\]](#)

Weight: 1 octet carrying the weight used for load-balancing purposes. The use of weight is described in [Section 3.4 of \[RFC8402\]](#).

Algorithm: Refer to ISIS/OSPFv2/OSPFv3 LAN-Adj-SID per Algorithm([\[I-D.ietf-lsr-algorithm-related-adjacency-sid\]](#)).

Reserved: 1-octet that MUST be set to 0 and ignored on receipt.

Neighbor ID/IS-IS System ID: Refer to ISIS/OSPFv2/OSPFv3 LAN-Adj-SID per Algorithm([\[I-D.ietf-lsr-algorithm-related-adjacency-sid\]](#)).

SID/Label/Index: Refer to ISIS/OSPFv2/OSPFv3 LAN-Adj-SID per Algorithm([\[I-D.ietf-lsr-algorithm-related-adjacency-sid\]](#)).

3.2. BGP-LS Extensions for Segment Descriptors

[I-D.peng-idr-segment-routing-te-policy-attr] defines four new Segment types of Segment Sub-TLV to provide algorithm information for SR-MPLS Adjacency-SIDs. This section defines the encoding of the Segment Descriptors for each of those Segment types to be used in the Segment sub-TLV describes in [I-D.ietf-idr-bgp-ls-sr-policy].

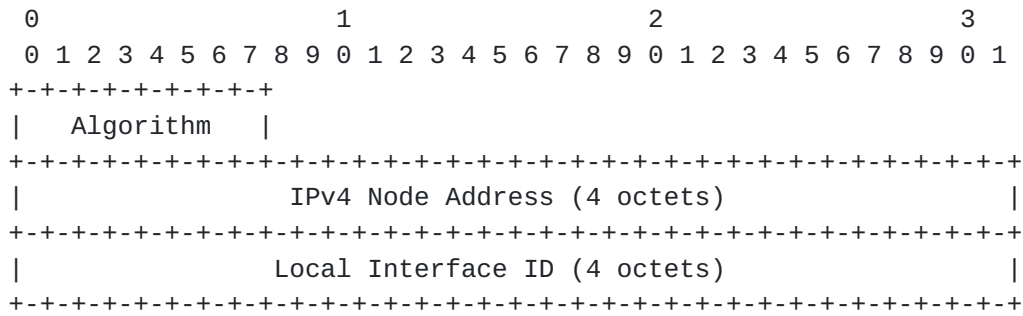
The following Segment Descriptors are defined and their mapping to the respective segment types defined in [I-D.peng-idr-segment-routing-te-policy-attr]:

Type	Segment Description
TBD3	(Type L)IPv4 Address and Local Interface ID with optional Algorithm
TBD4	(Type M)IPv4 Addresses for link endpoints as Local, Remote pair with optional Algorithm
TBD5	(Type N)IPv6 Prefix and Interface ID for link endpoints as Local, Remote pair, with optional Algorithm for SR-MPLS
TBD6	(Type O)IPv6 Addresses for link endpoints as Local, Remote pair, with optional Algorithm for SR-MPLS

Table 2: New Segment Descriptors

3.2.1. Type TBD3: IPv4 Address and Local Interface ID with optional Algorithm

The Segment is SR-MPLS Adjacency SID with optional Algorithm type and is specified as an IPv4 node address along with the local interface ID on that node. The format of its Segment Descriptor is as follows:



where:

Algorithm: 1-octet value that indicates the algorithm used for picking the SID. This is valid only when the A-flag has been set in the Segment TLV as defined in [[I-D.ietf-idr-bgp-ls-sr-policy](#)]. When A-Flag is not encoded, this field SHOULD be set to zero on transmission and MUST be ignored on receipt.

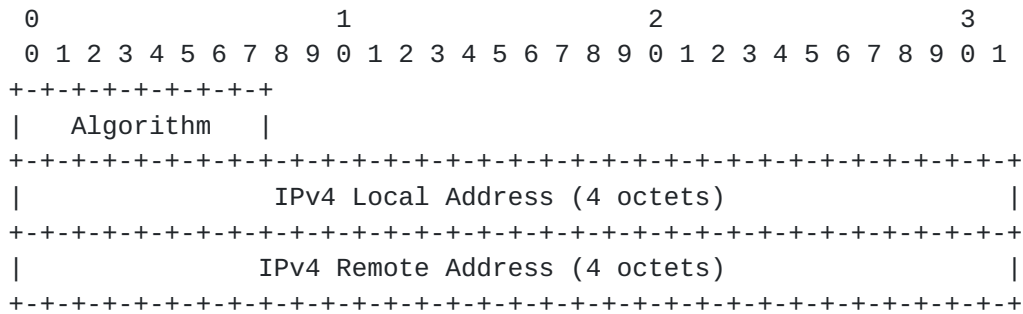
Other fields have the same meaning as the Type G Segment Sub-TLV Section 5.8.1.5 of [[I-D.ietf-idr-bgp-ls-sr-policy](#)], where:

IPv4 Node Address: 4-octet value which carries the IPv4 address associated with the node

Local Interface ID: 4-octet value which carries the local interface ID of the node identified by the Node Address.

3.2.2. Type TBD4: IPv4 Addresses for link endpoints as Local, Remote pair with optional Algorithm

The Segment is SR-MPLS Adjacency SID with optional Algorithm type and is specified as a pair of IPv4 local and remote addresses. The format of its Segment Descriptor is as follows:



where:

Algorithm: 1-octet value that indicates the algorithm used for picking the SID. This is valid only when the A-flag has been set in the Segment TLV as defined in [I-D.ietf-idr-bgp-ls-sr-policy]. When A-Flag is not encoded, this field SHOULD be set to zero on transmission and MUST be ignored on receipt.

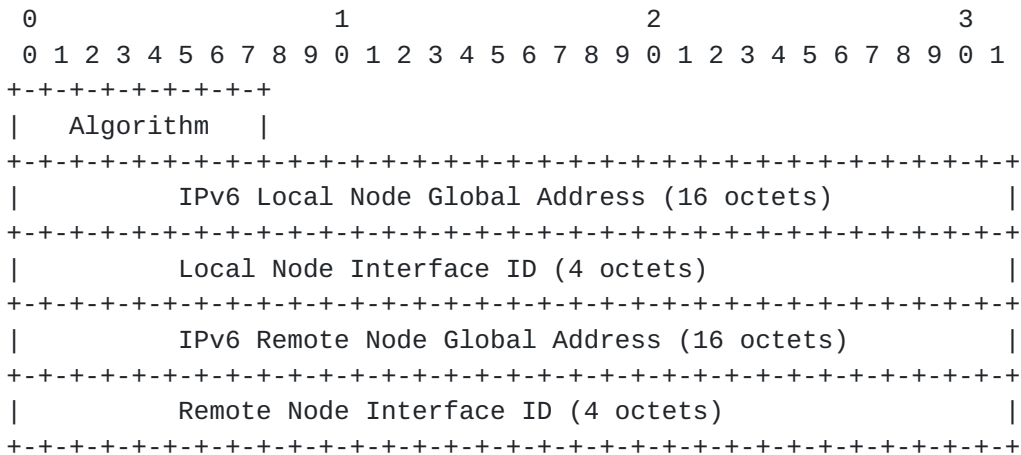
Other fields have the same meaning as the Type G Segment Sub-TLV Section 5.8.1.6 of [I-D.ietf-idr-bgp-ls-sr-policy], where:

IPv4 Local Address: 4-octet value which carries the local IPv4 address associated with the node.

IPv4 Remote Address: 4-octet value which carries the remote IPv4 address associated with the node's neighbor. This is optional and MAY be set to 0 when not used (e.g. when identifying point-to-point links).

3.2.3. Type TBD5: IPv6 Prefix and Interface ID for link endpoints as Local, Remote pair, with optional Algorithm for SR-MPLS

The Segment is SR-MPLS Adjacency SID optional Algorithm type and is specified as a pair of IPv6 global address and interface ID for local and remote nodes. The format of its Segment Descriptor is as follows:



where:

Algorithm: 1-octet value that indicates the algorithm used for picking the SID. This is valid only when the A-flag has been set in the Segment TLV as defined in [I-D.ietf-idr-bgp-ls-sr-policy]. When A-Flag is not encoded, this field SHOULD be set to zero on transmission and MUST be ignored on receipt.

Other fields have the same meaning as the Type G Segment Sub-TLV Section 5.8.1.7 of [[I-D.ietf-idr-bgp-ls-sr-policy](#)], where:

IPv6 Local Node Global Address: 16-octet value which carries the IPv6 global address associated with the local node.

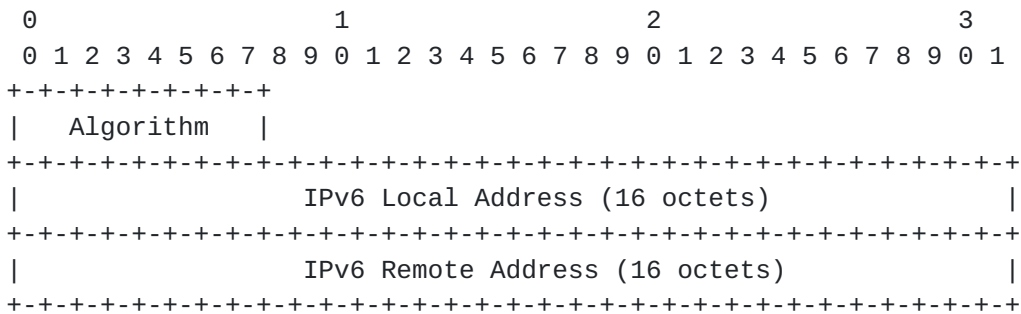
Local Node Interface ID : 4-octet value which carries the interface ID of the local node identified by the Local Node Address.

IPv6 Remote Node Global Address: 16-octet value which carries the IPv6 global address associated with the remote node. This is optional and MAY be set to 0 when not used (e.g. when identifying point-to-point links).

Remote Node Interface ID: 4-octet value which carries the interface ID of the remote node identified by the Remote Node Address. This is optional and MAY be set to 0 when not used (e.g. when identifying point-to-point links).

3.2.4. Type TBD6: IPv6 Addresses for link endpoints as Local, Remote pair, with optional Algorithm for SR-MPLS

The Segment is SR-MPLS Adjacency SID optional Algorithm type and is specified as a pair of IPv6 Global addresses for local and remote interface addresses. The format of its Segment Descriptor is as follows:



where:

Algorithm: 1-octet value that indicates the algorithm used for picking the SID. This is valid only when the A-flag has been set in the Segment TLV as defined in [[I-D.ietf-idr-bgp-ls-sr-policy](#)]. When A-Flag is not encoded, this field SHOULD be set to zero on transmission and MUST be ignored on receipt.

Other fields have the same meaning as the Type G Segment Sub-TLV Section 5.8.1.8 of [[I-D.ietf-idr-bgp-ls-sr-policy](#)], where:

IPv6 Local Address: 16-octet value which carries the local IPv6 address associated with the node.

IPv6 Remote Address: 16-octet value which carries the remote IPv6 address associated with the node's neighbor

4. Acknowledgements

TBD.

5. IANA Considerations

5.1. BGP-LS Link Attribute TLVs registry

This document makes the following Link Attribute TLVs registry under the "Border Gateway Protocol-Link State (BGP-LS) Parameter".

Value	Description	Reference
TBD1	Adjacency SID per Algorithm	This document
TBD2	LAN Adjacency SID per Algorithm	This document

5.2. BGP-LS SR Segment Descriptors

This document makes the following SR Segment Descriptor Types registry under the "SR Segment Descriptor Types". The registry contains the following codepoints to be assigned by IANA with the reference set to this document:

Value	Segment Description	Reference
TBD3	IPv4 Address and Local Interface ID with optional Algorithm	This document
TBD4	IPv4 Addresses for link endpoints as Local, Remote pair with optional Algorithm	This document
TBD5	IPv6 Prefix and Interface ID for link endpoints as Local, Remote pair, with optional Algorithm for SR-MPLS	This document
TBD6	IPv6 Addresses for link endpoints as Local, Remote pair, with optional Algorithm for SR-MPLS	This document

6. Security Considerations

TBD.

7. Normative References

[I-D.ietf-idr-bgp-ls-sr-policy]

Previdi, S., Talaulikar, K., Dong, J., Gredler, H., and J.
Tantsura, "Advertisement of Segment Routing Policies using

Chen, et al.

Expires 23 March 2024

[Page 10]

BGP Link-State", Work in Progress, Internet-Draft, [draft-ietf-idr-bgp-ls-sr-policy-01](https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-01), 23 July 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-01>>.

[I-D.ietf-lsr-algorithm-related-adjacency-sid]

Peng, S., Chen, R., Talaulikar, K., and P. Psenak, "Algorithm Related IGP-Adjacency SID Advertisement", Work in Progress, Internet-Draft, [draft-ietf-lsr-algorithm-related-adjacency-sid-05](https://datatracker.ietf.org/doc/html/draft-ietf-lsr-algorithm-related-adjacency-sid-05), 25 May 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-algorithm-related-adjacency-sid-05>>.

[I-D.peng-idr-segment-routing-te-policy-attr]

Liu, Y. and S. Peng, "Advertising SID Algorithm Information in BGP", Work in Progress, Internet-Draft, [draft-peng-idr-segment-routing-te-policy-attr-05](https://datatracker.ietf.org/doc/html/draft-peng-idr-segment-routing-te-policy-attr-05), 28 August 2023, <<https://datatracker.ietf.org/doc/html/draft-peng-idr-segment-routing-te-policy-attr-05>>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](https://www.rfc-editor.org/info/rfc2119), [RFC 2119](https://www.rfc-editor.org/info/rfc2119), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](https://www.rfc-editor.org/info/rfc2119) Key Words", [BCP 14](https://www.rfc-editor.org/info/rfc8174), [RFC 8174](https://www.rfc-editor.org/info/rfc8174), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8402] Filssils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](https://www.rfc-editor.org/info/rfc8402), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

[RFC9085] Previdi, S., Talaulikar, K., Ed., Filssils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", [RFC 9085](https://www.rfc-editor.org/info/rfc9085), DOI 10.17487/RFC9085, August 2021, <<https://www.rfc-editor.org/info/rfc9085>>.

Authors' Addresses

Ran Chen
ZTE Corporation
Nanjing
China
Email: chen.ran@zte.com.cn

Shaofu Peng
ZTE Corporation
Nanjing
China
Email: peng.shaofu@zte.com.cn

Peter Psenak
Cisco Systems
Slovakia
Email: ppsenak@cisco.com

Liu Yao
ZTE Corporation
China
Email: liu.yao71@zte.com.cn

