Workgroup: Network
Internet-Draft:
draft-peng-idr-segment-routing-te-policy-
attr-01
Published: 19 August 2021
Intended Status: Standards Track
Expires: 20 February 2022
Authors: L. Yao P. Shaofu
ZTE Corporation ZTE Corporation

Advertising Segment Routing Policies Attributes in BGP

Abstract

This document proposes extensions of BGP and defines some new Segment Types with algorithm information to meet more requirements when delivering SR Policy via BGP.

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 20 February 2022.

Copyright Notice

Table of Contents

1. Introduction
   1.1. Requirements Language
2. New Segment Types for SR-MPLS Adjacency with optional Algorithm
   2.1. Type M: IPv4 Address + Local Interface ID with optional Algorithm
   2.2. Type N: IPv4 Local and Remote addresses with optional Algorithm
   2.3. Type O: IPv6 Address + Interface ID for local and remote pair with optional Algorithm related SID for SR MPLS
   2.4. Type P: IPv6 Local and Remote addresses with optional Algorithm for SR MPLS
3. New Segment Types for SID only, with optional Algorithm
   3.1. Type L: MPLS SID only, with optional Algorithm
   3.2. Type Q: SRv6 SID only, with optional Algorithm
4. IANA Considerations
5. Security Considerations
6. Acknowledgements
7. Normative References

Authors' Addresses

1. Introduction

Segment Routing (SR) [RFC8402] allows a headend node to steer a packet flow along any path. Intermediate per-flow states are eliminated thanks to source routing.

[I-D.ietf-spring-segment-routing-policy] details the concepts of SR Policy and steering into an SR Policy. These apply equally to the MPLS and IPv6 data plane instantiations of Segment Routing with their respective representations of segments as SR-MPLS SID and SRv6 SID as described in [RFC8402].

[I-D.ietf-idr-segment-routing-te-policy] specifies the way to use BGP to distribute one or more of the candidate paths of an SR Policy to the headend of that policy. It defines a new BGP address family (SAFI), i.e., SR Policy SAFI NLRI. In UPDATE messages of that address family, the NLRI identifies an SR Policy Candidate Path, and the attributes encode the segment lists and other details of that SR Policy Candidate Path. 11 Segment Types (from A to K) are defined to encode SR-MPLS or SRv6 segments.

As specified in [I-D.ietf-idr-segment-routing-te-policy], the SR algorithm can be optionally specified for Segment Types C(IPv4 Node and SID), D(IPv6 Node and SID for SR-MPLS), I(IPv6 Node and SID for SRv6), J(IPv6 Node, index for remote and local pair, and SID for SRv6), and K(IPv6 Local/Remote addresses and SID for SRv6). That is, currently the algorithm can be carried along with SR-MPLS prefix

SID, SRv6 prefix SID and SRv6 adjacency SID when delivering SR Policy via BGP.

This document proposes extensions of BGP and defines some new Segment Types with algorithm information to meet more requirements when delivering SR Policy via BGP.

1.1. 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.

2. New Segment Types for SR-MPLS Adjacency with optional Algorithm

[I-D.ietf-lsr-algorithm-related-adjacency-sid] complements that the algorithm can be also included as part of an Adjacency-SID advertisement for SR-MPLS, in scenarios where multiple algorithm share the same link resource. In this case, an SR-MPLS Policy advertised to the headend may also contain algorithm specific Adjacency-SID.

This section defines 4 new Segment Sub-TLVs of Segment List Sub-TLV to provide algorithm information for SR-MPLS Adjacency-SID.

The processing procedures for SID with algorithm specified in [I-D.ietf-spring-segment-routing-policy] and [I-D.ietf-idr-segment-routing-te-policy] are still applicable for the new segment types. When the algorithm is not specified for the SID types above which optionally allow for it, the headend SHOULD use the Strict Shortest Path algorithm if available; otherwise, it SHOULD use the default Shortest Path algorithm.

2.1. Type M: IPv4 Address + Local Interface ID with optional Algorithm

The Type M Segment Sub-TLV is similar with existed Type E Segment Sub-TLV, it also encodes an IPv4 node address, a local interface Identifier (Local Interface ID) and an optional SR-MPLS SID, but with additional algorithm information. The format is as follows:

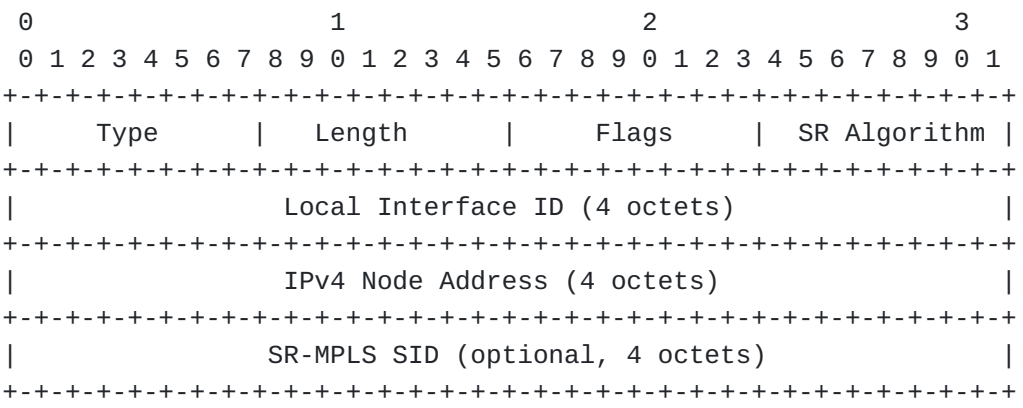

Figure 1

Where:

SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [RFC8402] when A-Flag as defined in section 2.4.4.2.12 [I-D.ietf-idr-segment-routing-te-policy] is present. SR Algorithm is used by SRPM as described in section 4 in [I-D.ietf-spring-segment-routing-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 existing Type E Segment Sub-TLV.

2.2. Type N: IPv4 Local and Remote addresses with optional Algorithm

The Type N Segment Sub-TLV is similar with existed Type F Segment Sub-TLV, it also encodes an adjacency local address, an adjacency remote address and an optional SR-MPLS SID, but with additional algorithm information. The format is as follows:

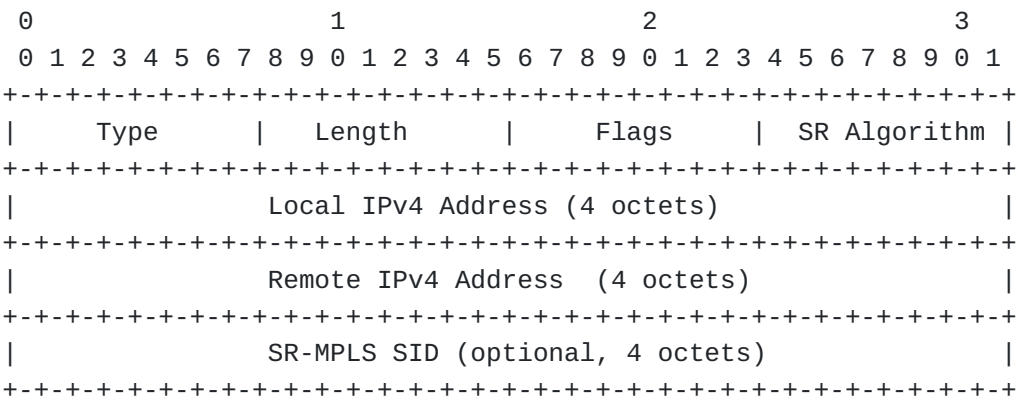

Figure 2

Where:

SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [RFC8402] when A-Flag as defined in section 2.4.4.2.12 [I-D.ietf-idr-segment-routing-te-policy] is present. SR Algorithm is used by SRPM as described in section 4 in [I-D.ietf-spring-segment-routing-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 existed Type F Segment Sub-TLV.

2.3. Type O: IPv6 Address + Interface ID for local and remote pair with optional Algorithm related SID for SR MPLS

The Type O Segment Sub-TLV is similar with existed Type G Segment Sub-TLV, it also encodes an IPv6 Link Local adjacency with IPv6 local node address, a local interface identifier (Local Interface ID), IPv6 remote node address , a remote interface identifier (Remote Interface ID) and an optional SR-MPLS SID, but with additional algorithm information. The format is as follows:

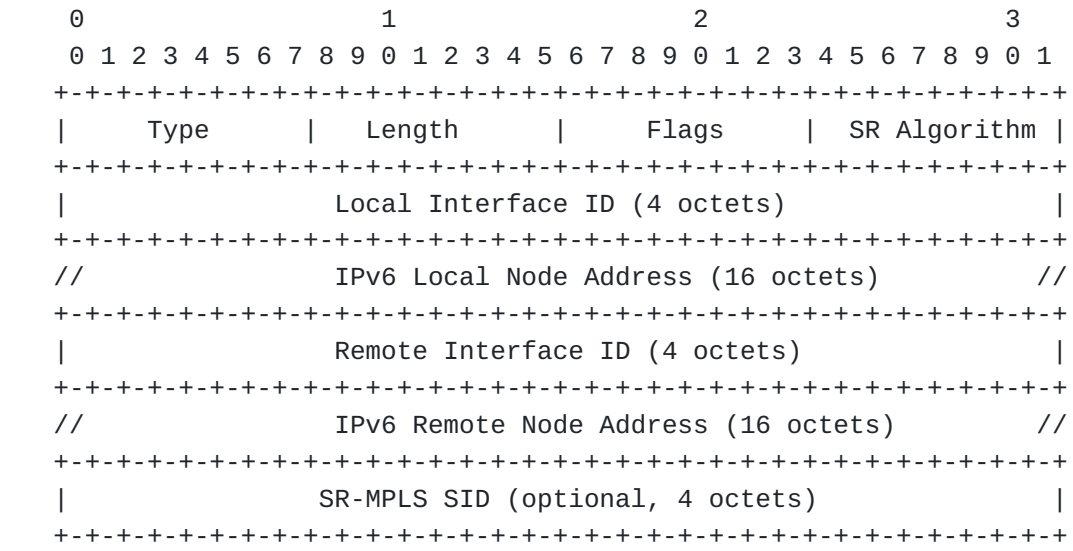

Figure 3

Where:

SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [RFC8402] when A-Flag as defined in section 2.4.4.2.12 [I-D.ietf-idr-segment-routing-te-policy] is present. SR Algorithm is used by SRPM as described in section 4 in [I-D.ietf-spring-segment-routing-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 existed Type G Segment Sub-TLV.

2.4. Type P: IPv6 Local and Remote addresses with optional Algorithm for SR MPLS

The Type P Segment Sub-TLV is similar with existed Type H Segment Sub-TLV, it also encodes an adjacency local address, an adjacency remote address and an optional SR-MPLS SID, but with additional algorithm information. The format is as follows:

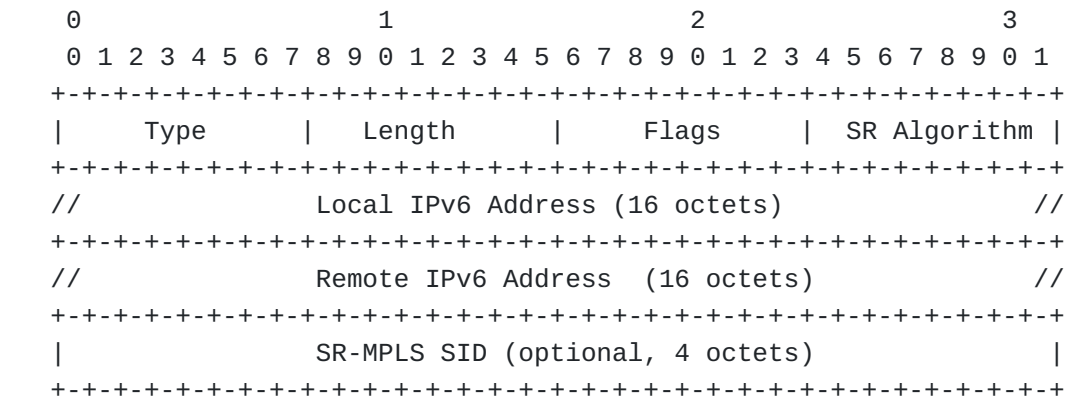

Figure 4

Where:

SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [RFC8402] when A-Flag as defined in section 2.4.4.2.12 [I-D.ietf-idr-segment-routing-te-policy] is present. SR Algorithm is used by SRPM as described in section 4 in [I-D.ietf-spring-segment-routing-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 existed Type H Segment Sub-TLV.

3. New Segment Types for SID only, with optional Algorithm

Segment Sub-TLV for Type A defined in section 2.4.4.2.1 [I-D.ietf-idr-segment-routing-te-policy] carries only the SID information in the form of MPLS Label. Segment Sub-TLV for Type B defined in section 2.4.4.2.2 [I-D.ietf-idr-segment-routing-te-policy] carries only the SID information in the form of IPv6 address.

If the algorithm information is carried along with the SIDs, it's useful in the scenarios below:

Scenario 1: The algorithm may be optionally provided to the headend for verification purposes. The headend can check if the SID value and the related algorithm received can be found in its SR-DB if requested to do so.

Scenario 2: The headend may not know about the SID-related algorithm especially in the inter-domain scenario. Providing the algorithm information benefits troubleshooting and network management.

This section defines 2 new Segment Sub-TLVs of Segment List Sub-TLV to provide algorithm information for SR-MPLS/SRv6 SID.

3.1. Type L: MPLS SID only, with optional Algorithm

The Type L Segment Sub-TLV is similar with the Type A Segment Sub-TLV, it also encodes a single SR-MPLS SID, but with additional algorithm information. The format is as follows:

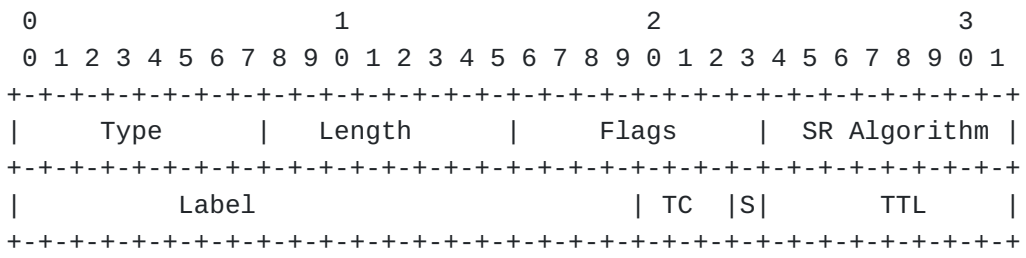

Figure 5

Where:

SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [RFC8402] when A-Flag as defined in section 2.4.4.2.12 [I-D.ietf-idr-segment-routing-te-policy] is present. 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 Type A Segment Sub-TLV.

3.2. Type Q: SRv6 SID only, with optional Algorithm

The Type Q Segment Sub-TLV is similar with existed Type B Segment Sub-TLV, it also encodes a single SRv6 SID, but with additional algorithm, endpoint behavior and SID strucutre information. The format is as follows:

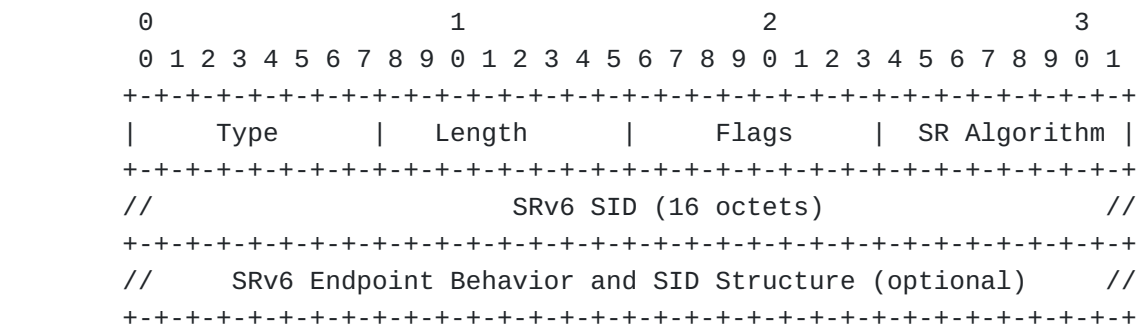

Figure 6

Where:

Type: TBD

Length is variable.

SR Algorithm: 1 octet specifying SR Algorithm as described in section 3.1.1 in [RFC8402] when A-Flag as defined in section 2.4.4.2.12 [I-D.ietf-idr-segment-routing-te-policy] is present. 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 B Segment Sub-TLV.

4. IANA Considerations

This document requests codepoint allocations for new Segment Sub-TLVs in the "SR Policy List Sub-TLVs" registry.

Value	Description	Reference
TBD1	Type L MPLS Algorithm related SID sub-TLV	This document
TBD2	Type M IPv4 Node, index and Algorithm related SID sub-TLV	This document
TBD3	Type N IPv4 Local/Remote addresses and Algorithm related SID sub-TLV	This document
TBD4	Type O IPv6 Node, index for remote and local pair and Algorithm related SID for SR-MPLS sub-TLV	This document
TBD5	Type P IPv6 Local/Remote addresses and Algorithm related SID sub-TLV	This document
TBD6	Type Q SRv6 Algorithm related SID sub-TLV	This document

Figure 7

5. Security Considerations

Procedures and protocol extensions defined in this document do not affect the security considerations discussed in [I-D.ietf-idr-segment-routing-te-policy].

6. Acknowledgements

TBD

7. Normative References

[I-D.ietf-idr-bgp-ls-segment-routing-ext]
Previdi, S., Talaulikar, K., Filsfils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-segment-routing-ext-18, 15 April 2021, <https://www.ietf.org/archive/id/draft-ietf-idr-bgp-ls-segment-routing-ext-18.txt>.

[I-D.ietf-idr-bgpls-srv6-ext]
Dawra, G., Filsfils, C., Talaulikar, K., Chen, M., Bernier, D., and B. Decraene, "BGP Link State Extensions for SRv6", Work in Progress, Internet-Draft, draft-ietf-idr-bgpls-srv6-ext-08, 8 June 2021, <https://www.ietf.org/archive/id/draft-ietf-idr-bgpls-srv6-ext-08.txt>.

[I-D.ietf-idr-segment-routing-te-policy]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Rosen, E., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-segment-routing-te-policy-13, 7 June 2021, <https://www.ietf.org/archive/id/draft-ietf-idr-segment-routing-te-policy-13.txt>.

[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-00, 10 June 2021, <https://www.ietf.org/archive/id/draft-ietf-lsr-algorithm-related-adjacency-sid-00.txt>.

[I-D.ietf-lsr-flex-algo] Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", Work in Progress, Internet-Draft, draft-ietf-lsr-flex-algo-17, 6 July 2021, <https://www.ietf.org/archive/id/draft-ietf-lsr-flex-algo-17.txt>.

[I-D.ietf-lsr-isis-srv6-extensions]
Psenak, P., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over IPv6 Dataplane", Work in Progress, Internet-Draft, draft-ietf-lsr-isis-srv6-extensions-17, 18 June 2021, <https://www.ietf.org/archive/id/draft-ietf-lsr-isis-srv6-extensions-17.txt>.

[I-D.ietf-lsr-ospfv3-srv6-extensions] Li, Z., Hu, Z., Cheng, D., Talaulikar, K., and P. Psenak, "OSPFv3 Extensions for SRv6", Work in Progress, Internet-Draft, draft-ietf-lsr-ospfv3-srv6-extensions-02, 15 February 2021, <https://www.ietf.org/archive/id/draft-ietf-lsr-ospfv3-srv6-extensions-02.txt>.

[I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, draft-ietf-spring-segment-routing-policy-13, 28 May 2021, <https://www.ietf.org/archive/id/draft-ietf-spring-segment-routing-policy-13.txt>.

[I-D.ietf-spring-srv6-network-programming]
Filsfils, C., Garvia, P. C., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-network-programming-28, 29 December 2020, <https://www.ietf.org/archive/id/draft-ietf-spring-srv6-network-programming-28.txt>.

[I-D.peng-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-peng-lsr-algorithm-related-adjacency-sid-03, 17 May 2021, <https://www.ietf.org/archive/id/draft-peng-lsr-algorithm-related-adjacency-sid-03.txt>.

[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>.

[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>.

[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>.

[RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <https://www.rfc-editor.org/info/rfc8660>.

[RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <https://www.rfc-editor.org/info/rfc8665>.

[RFC8666] Psenak, P., Ed. and S. Previdi, Ed., "OSPFv3 Extensions for Segment Routing", RFC 8666, DOI 10.17487/RFC8666, December 2019, <https://www.rfc-editor.org/info/rfc8666>.

[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>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

Authors' Addresses

Liu Yao
ZTE Corporation
China

Email: liu.yao71@zte.com.cn

Peng Shaofu
ZTE Corporation
China

Email: peng.shaofu@zte.com.cn