Workgroup: IDR
Internet-Draft:
draft-peng-idr-segment-routing-te-policy-
attr-03
Published: 5 June 2022
Intended Status: Standards Track
Expires: 7 December 2022
Authors: Y. Liu S. Peng
ZTE ZTE

Advertising SID Algorithm Information 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 7 December 2022.

Table of Contents

1. Introduction

Segment Routing (SR) [RFC8402] allows a headend node to steer a packet flow along any path. [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.

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 [RFC2119].

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

[I-D.ietf-lsr-algorithm-related-adjacency-sid] complements that besides Prefix-SID, 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.

3.1. Type M: IPv4 Address and 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:

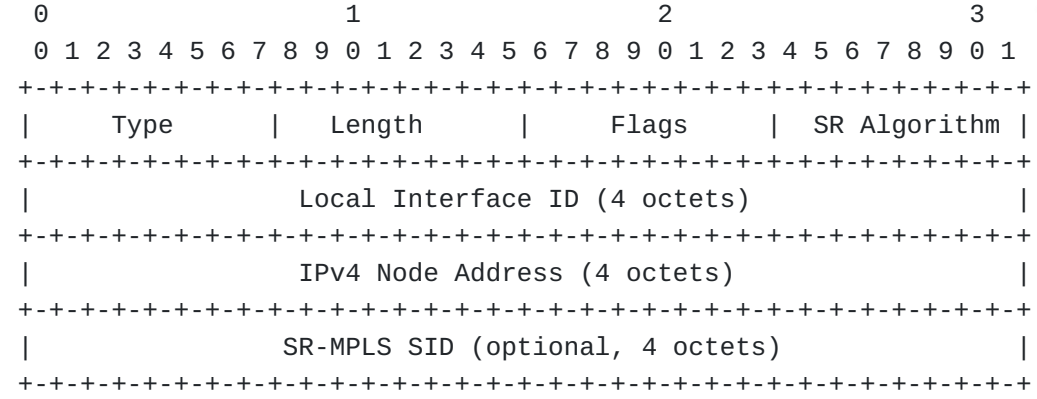

Where:

Type: TBD1

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.

3.2. Type N: IPv4 Addresses for link endpoints as Local, Remote pair 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:

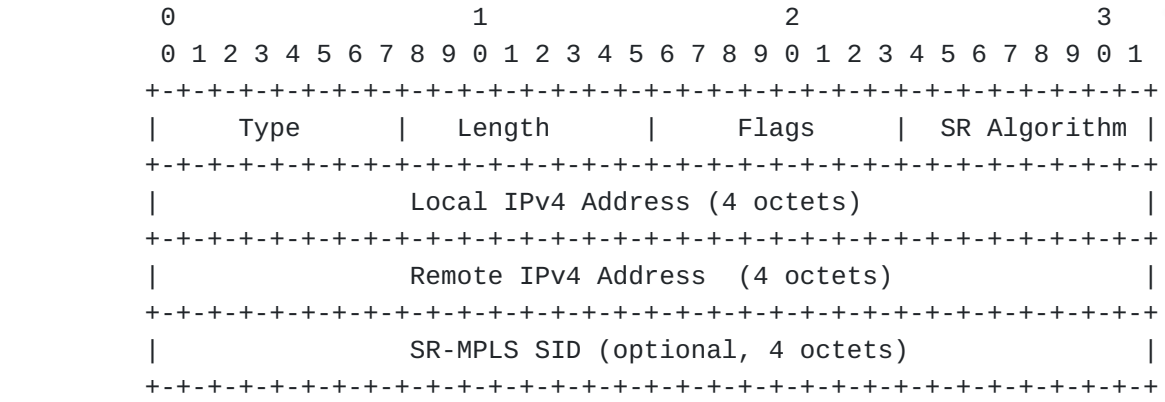

Where:

Type: TBD2

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.

3.3. Type O: IPv6 Prefix and Interface ID for link endpoints as Local, Remote pair, with optional Algorithm 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:

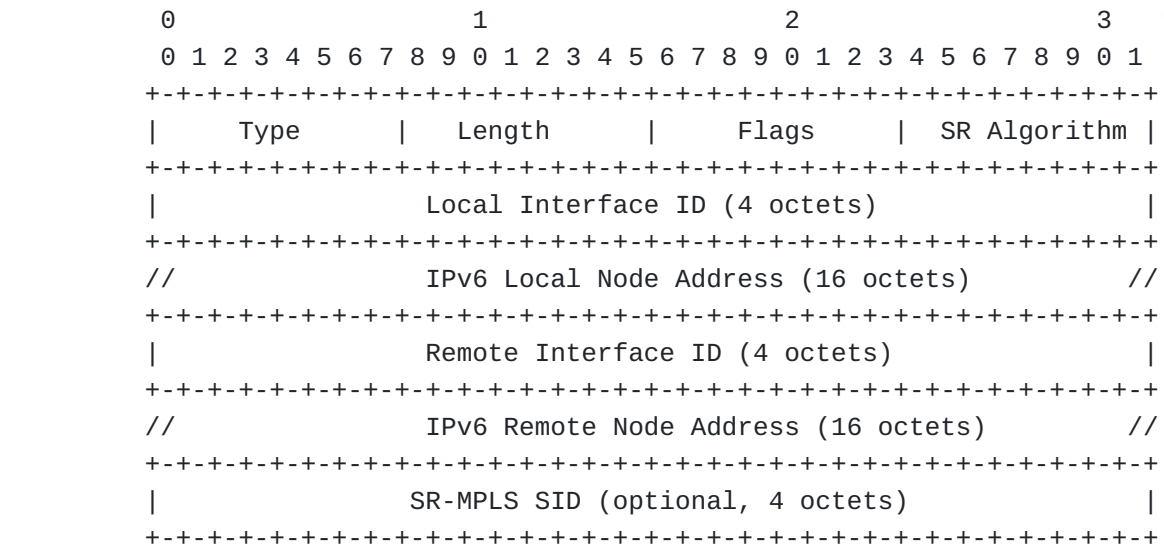

Where:

Type: TBD3

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.

3.4. Type P: IPv6 Addresses for link endpoints as Local, Remote pair, 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:

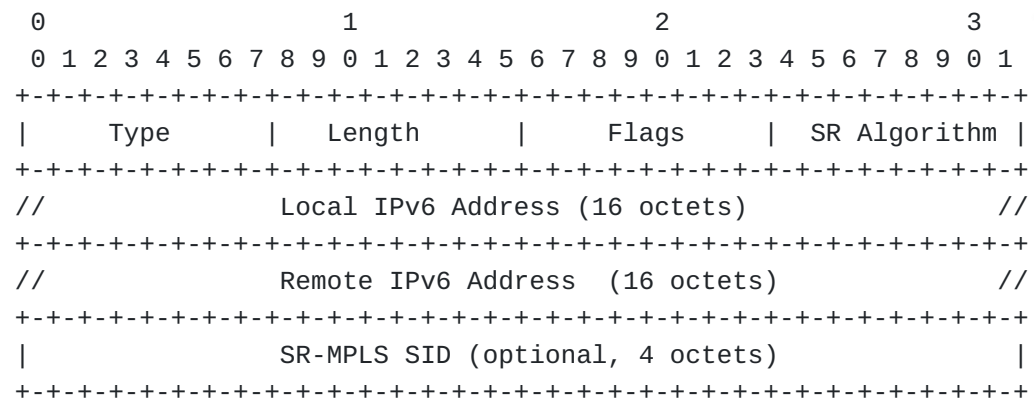

Where:

Type: TBD4

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.

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	Segment Type M sub-TLV	This document
TBD2	Segment Type N sub-TLV	This document
TBD3	Segment Type O sub-TLV	This document
TBD4	Segment Type P sub-TLV	This document

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. Acknowledgement

The authors would like to thank Ketan Talaulikar for his comments and suggestions.

7. References

7.1. Normative References

[I-D.ietf-idr-segment-routing-te-policy]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-segment-routing-te-policy-17, 14 April 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-segment-routing-te-policy-17>.

[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-22, 22 March 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-segment-routing-policy-22>.

[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>.

7.2. Informative References

[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-02, 18 January 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-lsr-algorithm-related-adjacency-sid-02>.

[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-20, 18 May 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-lsr-flex-algo-20>.

[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-18, 20 October 2021,

<https://datatracker.ietf.org/doc/html/draft-ietf-lsr-isis-srv6-extensions-18>.

[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-03, 19 November 2021, <https://datatracker.ietf.org/doc/html/draft-ietf-lsr-ospfv3-srv6-extensions-03>.

[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

Yao Liu

ZTE
Nanjing
China

Email: liu.yao71@zte.com.cn

Shaofu Peng
ZTE
Nanjing
China

Email: peng.shaofu@zte.com.cn