

Interdomain Routing Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 28, 2022

C. Li
Z. Li
Huawei Technologies
Y. Yin
China Telecom
W. Cheng
China Mobile
K. Talaulikar
Cisco Systems
January 24, 2022

SR Policy Extensions for Path Segment and Bidirectional Path
draft-ietf-idr-sr-policy-path-segment-05

Abstract

A Segment Routing (SR) policy is a set of candidate SR paths consisting of one or more segment lists with necessary path attributes. For each SR path, it may also have its own path attributes, and Path Segment is one of them. A Path Segment is defined to identify an SR path, which can be used for performance measurement, path correlation, and end-2-end path protection. Path Segment can be also used to correlate two unidirectional SR paths into a bidirectional SR path which is required in some scenarios, for example, mobile backhaul transport network.

This document defines extensions to BGP to distribute SR policies carrying Path Segment and bidirectional path information.

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](#)].

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

Internet-Draft Path ID and Bi-directional Path in BGP

January 2022

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 July 28, 2022.

Copyright Notice

Copyright (c) 2022 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	3
2.1.	Requirements Language	3
3.	Path Segment in SR Policy	3
3.1.	SR Path Segment Sub-TLV	5
4.	SR Policy for Bidirectional Path	7
4.1.	Reverse Path Segment List Sub-TLV	8
5.	Operations	8
6.	IANA Considerations	9
6.1.	Existing Registry: BGP Tunnel Encapsulation Attribute sub-TLVs	9
7.	Security Considerations	9
8.	Contributors	9
9.	Acknowledgements	10
10.	References	10
10.1.	Normative References	10
10.2.	Informative References	11
	Authors' Addresses	12

[1.](#) Introduction

Segment routing (SR) [[RFC8402](#)] is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. The ingress node steers packets into a specific path according to the Segment Routing Policy (SR Policy) as defined in [[I-D.ietf-spring-segment-routing-policy](#)]. For distributing SR

Internet-Draft Path ID and Bi-directional Path in BGP January 2022

policies to the headend, [[I-D.ietf-idr-segment-routing-te-policy](#)] specifies a mechanism by using BGP, and new sub-TLVs are defined for SR Policies in BGP UPDATE message.

In many use cases such as performance measurement, the path to which the packets belong is required to be identified. Furthermore, in some scenarios, for example, mobile backhaul transport network, there are requirements to support bidirectional path. However, there is no path identification information for each Segment List in the SR Policies defined in [[I-D.ietf-spring-segment-routing-policy](#)]. Also, the SR Policies defined in [[I-D.ietf-spring-segment-routing-policy](#)] only supports unidirectional SR paths.

Therefore, this document defines the extension to SR policies that carry Path Segment in the Segment List and support bidirectional path. The Path Segment can be a Path Segment in SR-MPLS [[I-D.ietf-spring-mpls-path-segment](#)] and SRv6 [[I-D.ietf-spring-srv6-path-segment](#)], or other IDs that can identify a path. Also, this document defines extensions to BGP to distribute SR policies carrying Path Segment and bidirectional path information.

[2.](#) Terminology

This memo makes use of the terms defined in [[RFC8402](#)] and [[I-D.ietf-idr-segment-routing-te-policy](#)].

[2.1.](#) 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.

[3.](#) Path Segment in SR Policy

As defined in [[I-D.ietf-idr-segment-routing-te-policy](#)] , the SR Policy encoding structure is as follows:

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>

Attributes:

 Tunnel Encaps Attribute (23)

 Tunnel Type: SR Policy

 Binding SID

 Preference

 Priority

 Policy Name

 Explicit NULL Label Policy (ENLP)

 Segment List

 Weight

 Segment

 Segment

 ...

 ...

An SR path can be specified by an Segment List sub-TLV that contains a set of segment sub-TLVs and other sub-TLVs as shown above. As defined in [[I-D.ietf-spring-segment-routing-policy](#)], a candidate path includes multiple SR paths specified by SID list. The Path Segment can be used for identifying an SR path(specified by SID list) from the headend and the tailend. Also, it can be used for identifying an SR candidate path in some use cases if needed. This document defines a new Path Segment sub-TLV within Segment List sub-TLV, the details will be described at [section 3.1](#). The new SR Policy encoding structure with Path Segment sub-TLV is expressed as below:

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>

Attributes:

 Tunnel Encaps Attribute (23)

 Tunnel Type: SR Policy

 Binding SID

 Preference

 Priority

 Policy Name

 Explicit NULL Label Policy (ENLP)

 Segment List

 Weight

 Path Segment

 Segment

 Segment

 ...

 Segment List

 Weight

 Path Segment

 Segment

 Segment

 ...

 ...

The Path Segment is used to identified an SR path, and it can be used in OAM or IOAM use cases. When all the SID Lists within a candidate path share the same Path Segment ID, the Path Segment can be used to collect the aggregated information of the candidate path. Multiple Path Segment MAY be included in a Segment List for different use cases, all of them SHOULD be inserted into the SID List.

3.1. SR Path Segment Sub-TLV

This section defines an SR Path Segment sub-TLV.

An SR Path Segment sub-TLV is included in the segment list sub-TLV to identify an SID list. It has the following format:

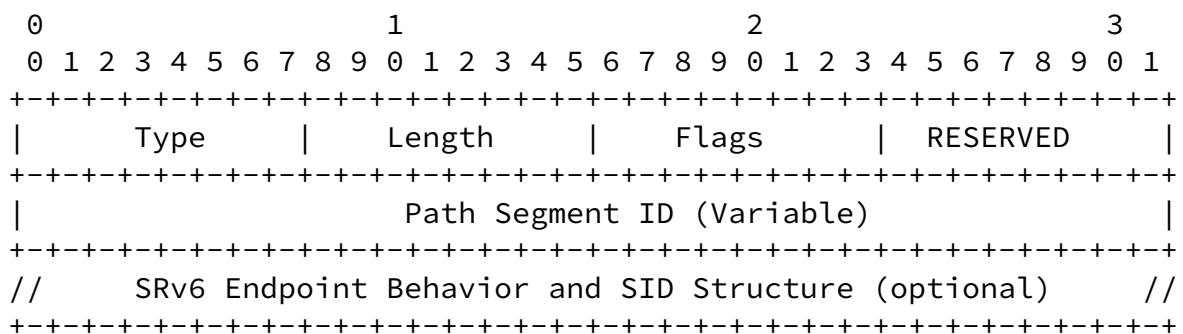


Figure 1. Path Segment sub-TLV

Where:

- o Type: to be assigned by IANA.

- o Length: the total length of the value field not including Type and Length fields.
- o Flags: 8 bits of flags. Following flags are defined:

```

0  1  2  3  4  5  6  7
+--+--+--+--+--+--+--+
|      Reserved      |B|L|
+--+--+--+--+--+--+--+

```

- o
 - * L-Flag: Local flag. Set when the Path Segment has local significance on an SR node.
 - * B-Flag: This flag, when set, indicates the presence of the SRv6 Endpoint Behavior and SID Structure encoding specified in Section 2.4.4.2.13 of [[I-D.ietf-idr-segment-routing-te-policy](#)]. It MUST be ignored when the value of length field is smaller than 18.
 - * The rest bits of Flag are reserved and MUST be set to 0 on transmission and MUST be ignored on receipt.
- o Path Segment ID: if the length is 2, then no Path Segment ID is present. If the length is 6 then the Path Segment ID is encoded in 4 octets [[I-D.ietf-spring-mpls-path-segment](#)] using the format below. TC, S, TTL (Total of 12 bits) are RESERVED and SHOULD be set to zero and MUST be ignored.

```

0                                1                                2                                3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|      Type      |      Length      |      Flags      |      Reserved      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|      Path Segment Label      | TC |S|      TTL      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 2. SR-MPLS Path Segment sub-TLV

If the length is 18 then the Path Segment ID contains a 16-octet SRv6 Path Segment ID [[I-D.ietf-spring-srv6-path-segment](#)].

If the length is larger than 18 and B-flag is set, then SRv6 Endpoint Behavior and SID Structure TLVs [[I-D.ietf-idr-segment-routing-te-policy](#)] is included.

[4.](#) SR Policy for Bidirectional Path

In some scenarios, for example, mobile backhaul transport network, there are requirements to support bidirectional path. In SR, a bidirectional path can be represented as a binding of two unidirectional SR paths. This document also defines a Reverse Segment List sub-TLV to describe the reverse path associated with the forward path specified by the Segment List. An SR policy carrying SR bidirectional path information is expressed as below:

```
SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>
  Attributes: Tunnel Encaps Attribute (23)
  Tunnel Type: SR Policy
    Binding SID
    Preference
    Priority
    Policy Name
    Explicit NULL Label Policy (ENLP)
    Segment List
      Weight
      Path Segment
      Segment
      Segment
      ...
    Reverse Segment List
      Path Segment
      Segment
      Segment
      ...
```

[4.1.](#) Reverse Path Segment List Sub-TLV

A Reverse Path Segment List sub-TLV is defined to specify an SR reverse path associated with the path specified by the Segment List, and it has the following format:

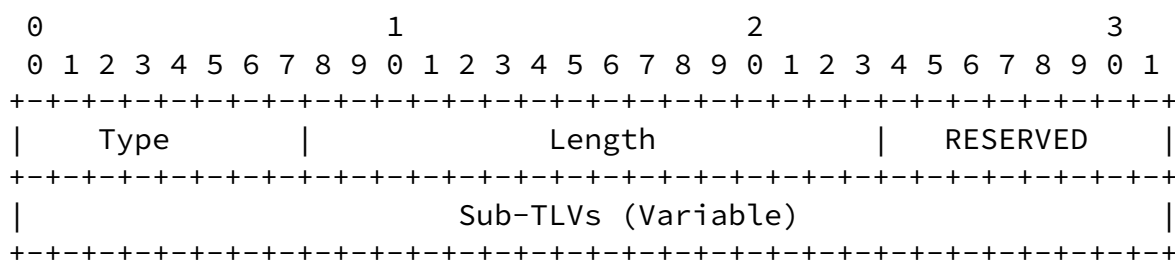


Figure 3. SR Reverse Path Segment List Sub-TLV

where:

Type: TBA.

Length: the total length of the sub-TLVs encoded within the Reverse Path Segment List Sub-TLV not including the Type and Length fields.

RESERVED: 1 octet of reserved bits. SHOULD be unset on transmission and MUST be ignored on receipt.

sub-TLVs, reuse the sub-TLVs in Segment List defined in [\[I-D.ietf-idr-segment-routing-te-policy\]](#).

- o One or more mandatory SR Path Segment sub-TLVs that contains the Path Segments of the reverse SR path.
- o One or more Segment sub-TLVs to specify the reverse SR path.

The Segment sub-TLVs in the Reverse Path Segment List sub-TLV provides the information of the reverse SR path, which can be used for directing egress BFD peer to use specific path for the reverse direction of the BFD session [\[I-D.ietf-mpls-bfd-directed\]](#) or other applications.

5. Operations

The document does not bring new operation beyond the description of operations defined in [\[I-D.ietf-idr-segment-routing-te-policy\]](#). The existing operations defined in [\[I-D.ietf-idr-segment-routing-te-policy\]](#) can apply to this document directly.

Typically but not limit to, the unidirectional or bidirectional SR policies carrying path identification information are configured by a controller.

After configuration, the unidirectional or bidirectional SR policies carrying path identification information will be advertised by BGP update messages. The operation of advertisement is the same as defined in [[I-D.ietf-idr-segment-routing-te-policy](#)], as well as the reception.

The consumer of the unidirectional or bidirectional SR policies is not the BGP process, it can be any applications, such as performance measurement [[I-D.gandhi-spring-udp-pm](#)]. The operation of sending information to consumers is out of scope of this document.

[6.](#) IANA Considerations

This document defines new Sub-TLVs in following registries:

[6.1.](#) Existing Registry: BGP Tunnel Encapsulation Attribute sub-TLVs

This document defines new sub-TLVs in the registry "SR Policy List Sub-TLVs" [[I-D.ietf-idr-segment-routing-te-policy](#)] to be assigned by IANA:

Codepoint	Description	Reference

TBA	Path Segment sub-TLV	This document
TBA	Reverse Segment List sub-TLV	This document

[7.](#) Security Considerations

TBA

[8.](#) Contributors

Mach(Guoyi) Chen
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: Mach.chen@huawei.com

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: jie.dong@huawei.com

James N Guichard
Futurewei Technologies
2330 Central Express Way
Santa Clara
USA

Email: james.n.guichard@futurewei.com

Huanan Chen
China Telecom
109 West Zhongshan Ave
Guangzhou
China

Email: chenhuan6@chinatelecom.cn

[9.](#) Acknowledgements

Many thanks to Shraddha Hedge for her detailed review and professional comments.

[10.](#) References

[10.1.](#) Normative References

Li, et al.

Expires July 28, 2022

[Page 10]

Internet-Draft Path ID and Bi-directional Path in BGP January 2022

[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", [draft-ietf-idr-segment-routing-te-policy-14](#) (work in progress), November 2021.

[I-D.ietf-spring-mpls-path-segment]

Cheng, W., Li, H., Chen, M., Gandhi, R., and R. Zigler, "Path Segment in MPLS Based Segment Routing Network", [draft-ietf-spring-mpls-path-segment-07](#) (work in progress), December 2021.

[I-D.ietf-spring-segment-routing-policy]

Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-14](#) (work in progress), October 2021.

[I-D.ietf-spring-srv6-path-segment]

Li, C., Cheng, W., Chen, M., Dhody, D., and Y. Zhu, "Path Segment for SRv6 (Segment Routing in IPv6)", [draft-ietf-spring-srv6-path-segment-03](#) (work in progress), November 2021.

[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>>.

[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>>.

[10.2.](#) Informative References

[I-D.gandhi-spring-udp-pm]
Gandhi, R., Filsfils, C., Voyer, D., Salsano, S., Ventre, P. L., and M. Chen, "UDP Path for In-band Performance Measurement for Segment Routing Networks", [draft-gandhi-spring-udp-pm-02](#) (work in progress), September 2018.

Li, et al.

Expires July 28, 2022

[Page 11]

Internet-Draft Path ID and Bi-directional Path in BGP January 2022

[I-D.ietf-mpls-bfd-directed]
Mirsky, G., Tantsura, J., Varlashkin, I., and M. Chen, "Bidirectional Forwarding Detection (BFD) Directed Return Path for MPLS Label Switched Paths (LSPs)", [draft-ietf-mpls-bfd-directed-18](#) (work in progress), August 2021.

Authors' Addresses

Cheng Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: c.l@huawei.com

Zhenbin Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Yuanyang Yin
China Telecom
Guangzhou
China

Email: yinyuany@chinatelecom.cn

Weiqiang Cheng
China Mobile
Beijing
China

Email: chengweiqiang@chinamobile.com

Ketan Talaulikar
Cisco Systems

Email: ketant.ietf@gmail.com