

Interdomain Routing Working Group
Internet-Draft
Intended status: Standards Track
Expires: 18 October 2025

C. Li
Z. Li
Huawei Technologies
Y. Zhu
China Telecom
W. Cheng
China Mobile
K. Talaulikar
Cisco Systems
16 April 2025

SR Policies Extensions for Path Segment and Bidirectional Path in BGP-LS
draft-ietf-idr-bgp-ls-sr-policy-path-segment-09

Abstract

This document specifies the way of collecting configuration and states of SR policies carrying Path Segment and bidirectional path information by using BGP-LS. Such information can be used by external components for many use cases such as performance measurement, path re-optimization and end-to-end protection.

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 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 18 October 2025.

Copyright Notice

Copyright (c) 2025 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. Terminology	3
3. Carrying SR Path Sub-TLVs in BGP-LS	4
3.1. SR Path Segment Sub-TLV	5
3.2. Reverse Segment List Sub-TLV	7
4. Operations	8
5. Error Handling and Fault Management	8
6. IANA Considerations	9
6.1. BGP-LS TLVs	9
7. Security Considerations	9
8. Contributors	10
9. Acknowledgements	10
10. References	10
10.1. Normative References	10
10.2. Informative References	12
Authors' Addresses	12

1. Introduction

Segment routing (SR) [RFC8402] is a source routing paradigm that allows the ingress node steers packets into a specific path according to the Segment Routing Policy [RFC9256].

However, the SR Policies defined in [RFC9256] only supports unidirectional SR paths and there is no path ID in a Segment List to identify an SR path. For identifying an SR path and supporting bidirectional path [RFC9545], the Path Segment and Reverse Path Segment List Sub-TLVs are defined for the Tunnel Encapsulation Attribute [RFC9012] for the SR Policy tunnel in [I-D.ietf-idr-sr-policy-path-segment]. The Path Segment identifier can be a Path Segment in SR-MPLS [RFC9545] and SRv6 [I-D.ietf-spring-srv6-path-segment], or other IDs that can identify the SR path.

In many network scenarios, the configuration and state of each TE Policy is required by a controller which allows the network operator to optimize several functions and operations through the use of a controller aware of both topology and state information [I-D.ietf-idr-bgp-ls-sr-policy].

To collect the TE Policy information that is locally available in a router, [I-D.ietf-idr-bgp-ls-sr-policy] describes a new mechanism by using BGP-LS update messages.

Based on the mechanism defined in [I-D.ietf-idr-bgp-ls-sr-policy], this document describes a mechanism to distribute configuration and states of the new SR policies defined in [I-D.ietf-idr-sr-policy-path-segment] to external components using BGP-LS.

2. Terminology

This document makes use of the terms defined in [RFC8402] and [I-D.ietf-idr-bgp-ls-sr-policy]. Some existing and new terms are listed below for reference.

- * SR: Segment Routing.
- * SR-MPLS: Segment Routing over MPLS data plane.
- * SRv6: Segment Routing over IPv6 data plane.
- * PSID: Path Segment Identifier.
- * SRPM: SR Policy Module [I-D.ietf-idr-sr-policy-safi].

3. Carrying SR Path Sub-TLVs in BGP-LS

A mechanism to collect states of SR Policies via BGP-LS is proposed by [I-D.ietf-idr-bgp-ls-sr-policy]. The characteristics of an SR policy can be described by a TE Policy State TLV, which is carried in the optional non-transitive BGP Attribute "LINK_STATE Attribute" defined in [RFC9552]. The TE Policy State TLV contains several sub-TLVs such as SR TE Policy sub-TLVs.

[I-D.ietf-idr-sr-policy-path-segment] defines the BGP extensions for Path Segment. The encoding is shown 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
 - ...
 - ...

Figure 1. Path Segment in SR policy

Also, [I-D.ietf-idr-sr-policy-path-segment] defines SR policy extensions for bidirectional SR path, the encoding is shown 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
      Weight
      Path Segment
      Segment
      Segment
      ...
```

Figure 2. SR policy for Bidirectional path

In order to collect configuration and states of unidirectional and bidirectional SR policies defined in [I-D.ietf-idr-sr-policy-path-segment], this document defines new sub-TLVs in SR TE Policy sub-TLVs.

3.1. SR Path Segment Sub-TLV

This section defines the SR Path Segment sub-TLV to describe a Path Segment, and it can be included in the Segment List sub-TLV as defined in [I-D.ietf-idr-bgp-ls-sr-policy]. An SR Path Segment sub-TLV can be associated with an SR path specified by a Segment List sub-TLV. Multiple Path Segment MAY be included in a Segment List for different 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. The format of Path Segment TLV is shown below.

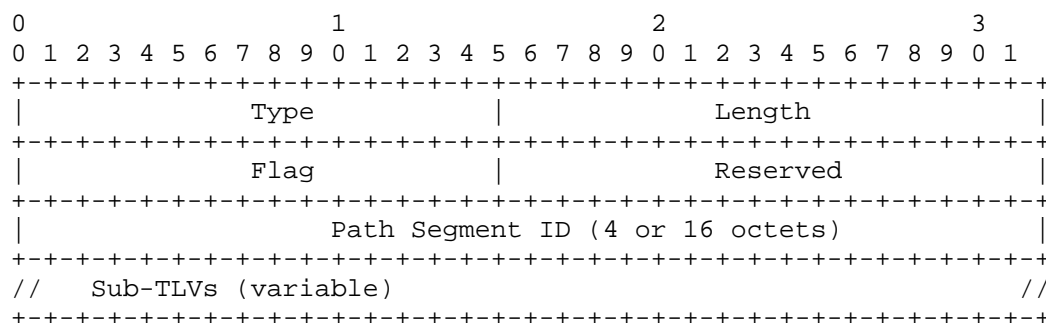



Figure 3. Path Segment sub-TLV

Where,

- * Type: to be assigned by IANA.
- * Length: the total length of the value field not including Type and Length fields.
- * Flags: 2 octet field that indicates attribute and status of the Path Segment. The following bit positions are defined. Other bits SHOULD be cleared by originator and MUST be ignored by receiver.



Where:

- * - D-Flag : Indicates the dataplane for the BSIDs. This flag is set when Path Segment ID is a 16-octet SRv6 SID. This flag is unset when the Path Segment ID is a 4-octet SR/MPLS label value.
- B-Flag: This flag when set indicates the presence of the SRv6 Endpoint Behavior and SID Structure encoding specified in [RFC9514]. The B-Flag when unset (clear) means that SRv6 Endpoint Behavior and SID Structure are not included. The B-Flag MUST be ignored when D-flag is unset. The B-Flag and D-Flag indicate the SRv6 Endpoint behavior and SID structure for the Path Segment ID value in the TLV.

- L-Flag: Local flag. Set when the Path Segment has local significance on an SR node. Unset when the Path Segment does not have local significance on an SR node
- * RESERVED: 2 octets. SHOULD be set to 0 by originator and MUST be ignored by receiver.
- * Path Segment ID: It indicates the Path Segment ID value based on the status flags.

The SRv6 Endpoint Behavior TLV (1250) and the SRv6 SID Structure TLV (1252) defined in [RFC9514] MAY be used as sub-TLVs of the SR Path Segment Sub-TLV. These optional sub-TLVs indicate the SRv6 Endpoint behavior and SID structure for the Path Segment ID value in the TLV when the Path Segment is an SRv6 Path Segment.

3.2. Reverse Segment List Sub-TLV

In some scenarios like 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 [RFC9545]. An SR policy carrying SR bidirectional path information is expressed in Figure 2. [I-D.ietf-idr-sr-policy-path-segment] defines a new sub-TLV to describe a reversed SR path of an SID list.

This section defines a Reverse Segment List sub-TLV to specify a reverse SR path associated with the path specified by the Segment List, and it reuses the format of SR Segment List TLV defined in [I-D.ietf-idr-bgp-ls-sr-policy]:




Figure 5. Reverse Segment List Sub-TLV

All fields, except the type are defined in [I-D.ietf-idr-bgp-ls-sr-policy], and this TLV reuses it directly. The Type of this TLV is TBA.

The SR Segment sub-TLV [I-D.ietf-idr-bgp-ls-sr-policy] MUST be included as an ordered set of sub-TLVs within the SR Segment List TLV when the SID-List is not empty. A SID-List may be empty in certain cases (e.g. for a dynamic path) where the headend has not yet performed the computation and hence not derived the segments required for the path; in such cases, the SR Segment List TLV SHOULD NOT include any SR Segment sub-TLVs [I-D.ietf-idr-bgp-ls-sr-policy]. In this case, the Path Segment Sub-TLV SHOULD NOT be included in the sub-TLVs field.

Note: currently, only one reverse SID list is supported, so the weight field CAN be ignored when processing. However, multiple reverse SID list MAY be supported in the future, and the use case of supporting this still need to be discussed.

4. Operations

The operations procedures of [RFC9552] can apply to this document.

Typically but not limited to, the uni/bidirectional SR policies carrying path identification information can be distributed by the ingress node.

Generally, BGP-LS is used for collecting link states and synchronizing with the external component. The consumer of the uni/bidirectional SR policies carrying path identification information is not BGP LS process by itself. This consumer can be any applications such as performance measurement [I-D.ietf-spring-stamp-srpm], path re-computation or re-optimization. The operation of sending information to other processes is out of scope of this document

5. Error Handling and Fault Management

This document defines a new SR Path Segment sub-TLV included in the Segment List sub-TLV as defined in [I-D.ietf-idr-bgp-ls-sr-policy], therefore, the error handling defined in [I-D.ietf-idr-bgp-ls-sr-policy] can apply to this document. The error handling as defined in [RFC7606] applies to new Sub-TLVs as well as SAFI context, therefore, the error handling in [RFC7606] also applies to this document.

Specifically, a BGP Speaker MUST perform Syntax validation of the SR Path Segment sub-TLV following the error handling defined in [RFC7606] and [I-D.ietf-idr-bgp-ls-sr-policy], to determine if it is malformed. This includes the validation of the length of the Sub-TLV and the range of the value fields. If any validation check fails, the Update message MUST be handle as 'Treat-as-withdraw'.

In addition, the validation of the individual fields of the TLVs/Sub-TLVs of the associated segment list are beyond the scope of BGP and out of the scope of this document. A BGP implementation MUST NOT perform semantic verification of such fields nor consider the SR Policy update to be invalid or not usable based on such validation. An implementation SHOULD log any errors found during the above validation for further analysis.

6. IANA Considerations

6.1. BGP-LS TLVs

IANA maintains a registry called "Border Gateway Protocol - Link State (BGP-LS) Parameters" with a sub-registry called "Node Anchor, Link Descriptor and Link Attribute TLVs". The following TLV codepoints are suggested (for early allocation by IANA):

Codepoint	Description	Reference
TBA	Path Segment sub-TLV	This document
TBA	Reverse Segment List sub-TLV	This document

7. Security Considerations

Similar to [I-D.ietf-idr-bgp-ls-sr-policy], the security mechanisms of the base BGP security model [RFC4271] apply to the extensions described in this document. Also, the new security considerations defined in [I-D.ietf-idr-bgp-ls-sr-policy] also apply to this document.

The Path Segment extension is included in the SR Policy extension [I-D.ietf-idr-bgp-ls-sr-policy], so it does not introduce extra security problems comparing the existing SR policy extension. The Path Segment information is critical to the path, and a wrong Path Segment ID may cause unexpected forwarding actions and results.

An implementation needs to make sure that the value of Path Segment ID is correct to avoid unexpected forwarding actions and results, especially in an SR-MPLS network. In addition, the Path Segment information distribution from a router to an controller has to be protected. The security considerations in [I-D.ietf-idr-bgp-ls-sr-policy] apply to this distribution procedure.

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

9. Acknowledgements

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

10. References

10.1. 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

BGP Link-State", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-sr-policy-17, 6 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-17>>.

[I-D.ietf-idr-sr-policy-safi]

Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-13, 6 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-13>>.

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

Li, C., Cheng, W., Chen, M., Dhody, D., and Y. Zhu, "Path Segment Identifier (PSID) in SRv6 (Segment Routing in IPv6)", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-path-segment-12, 3 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-path-segment-12>>.

[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>>.

[RFC7606] Chen, E., Ed., Scudder, J., Ed., Mohapatra, P., and K. Patel, "Revised Error Handling for BGP UPDATE Messages", RFC 7606, DOI 10.17487/RFC7606, August 2015, <<https://www.rfc-editor.org/info/rfc7606>>.

[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>>.

[RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.

- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.
- [RFC9514] Dawra, G., Filsfils, C., Talaulikar, K., Ed., Chen, M., Bernier, D., and B. Decraene, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing over IPv6 (SRv6)", RFC 9514, DOI 10.17487/RFC9514, December 2023, <<https://www.rfc-editor.org/info/rfc9514>>.
- [RFC9545] Cheng, W., Ed., Li, H., Li, C., Ed., Gandhi, R., and R. Zigler, "Path Segment Identifier in MPLS-Based Segment Routing Networks", RFC 9545, DOI 10.17487/RFC9545, February 2024, <<https://www.rfc-editor.org/info/rfc9545>>.
- [RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.

10.2. Informative References

- [I-D.ietf-idr-sr-policy-path-segment]
Li, C., Li, Z., Yin, Y., Cheng, W., and K. Talaulikar, "SR Policy Extensions for Path Segment and Bidirectional Path", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-path-segment-13, 2 October 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-path-segment-13>>.
- [I-D.ietf-spring-stamp-srpm]
Gandhi, R., Filsfils, C., Voyer, D., Chen, M., and R. F. Foote, "Performance Measurement Using Simple Two-Way Active Measurement Protocol (STAMP) for Segment Routing Networks", Work in Progress, Internet-Draft, draft-ietf-spring-stamp-srpm-17, 8 November 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-stamp-srpm-17>>.

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

Yongqing Zhu
China Telecom
109 West Zhongshan Ave
Guangzhou
China
Email: zhuyq8@chinatelecom.cn

Weiqiang Cheng
China Mobile
Beijing
China
Email: chengweiqiang@chinamobile.com

Ketan Talaulikar
Cisco Systems
Email: ketant.ietf@gmail.com