Interdomain Routing Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 30, 2022

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

January 26, 2022

SR Policies Extensions for Path Segment and Bidirectional Path in BGP-LS

draft-ietf-idr-bgp-ls-sr-policy-path-segment-02

Abstract

This document specifies the way of collecting configuration and states of SR policies carrying Path Segment and bidirectional path information by using BPG-LS. Such information can be used by external conponents 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 July 30, 2022.

Table of Contents

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

However, the SR Policies defined in [I-D.ietf-spring-segment-routing-policy] 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 [I-D.ietf-spring-mpls-path-segment], new policies carrying Path Segment and bidirectional path information are defined in [I-D.ietf-idr-sr-policy-path-segment], as well as the extensions to BGP to distribute new SR policies. 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.

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-te-lsp-distribution].

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

Based on the mechanism defined in [I-D.ietf-idr-te-lsp-distribution], 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 memo makes use of the terms defined in [RFC8402] and [I-D.ietf-idr-te-lsp-distribution].

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-te-lsp-distribution]. 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 [RFC7752]. 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-te-lsp-distribution] . 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.

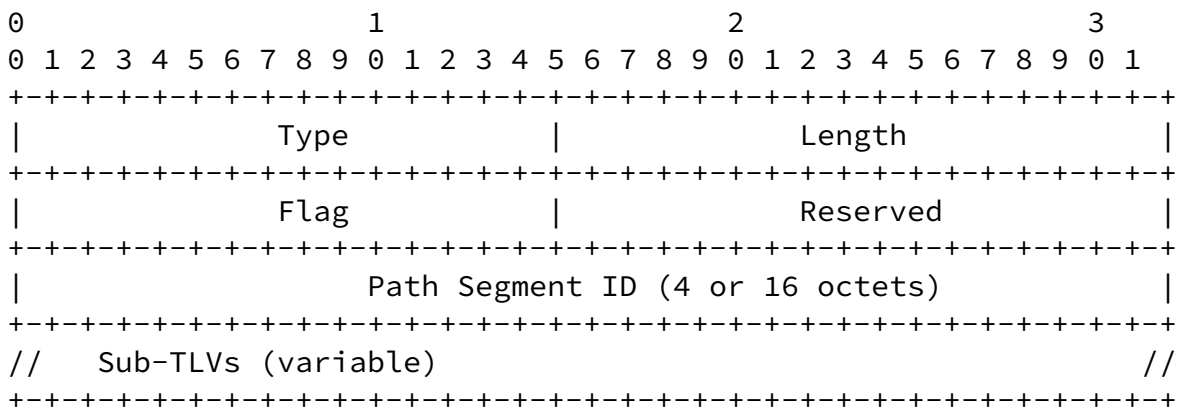

Figure 3. Path Segment sub-TLV

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.

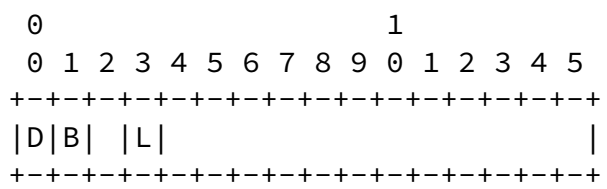

Where:

-

  - D-Flag : Indicates the dataplane for the BSIDs, it is set when Path Segment ID is a 16-octet SRv6 SID 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 [I-D.ietf-idr-bgpls-srv6-ext]. It MUST be ignored when D-flag is unset. They 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.

- 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 [I-D.ietf-idr-bgpls-srv6-ext] are used as sub-TLVs

of the SR Path Segment Sub-TLV to optionally indicate the SRv6 Endpoint behavior and SID structure for the Binding SID 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 [I-D.ietf-spring-mpls-path-segment]. 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-te-lsp-distribution]:

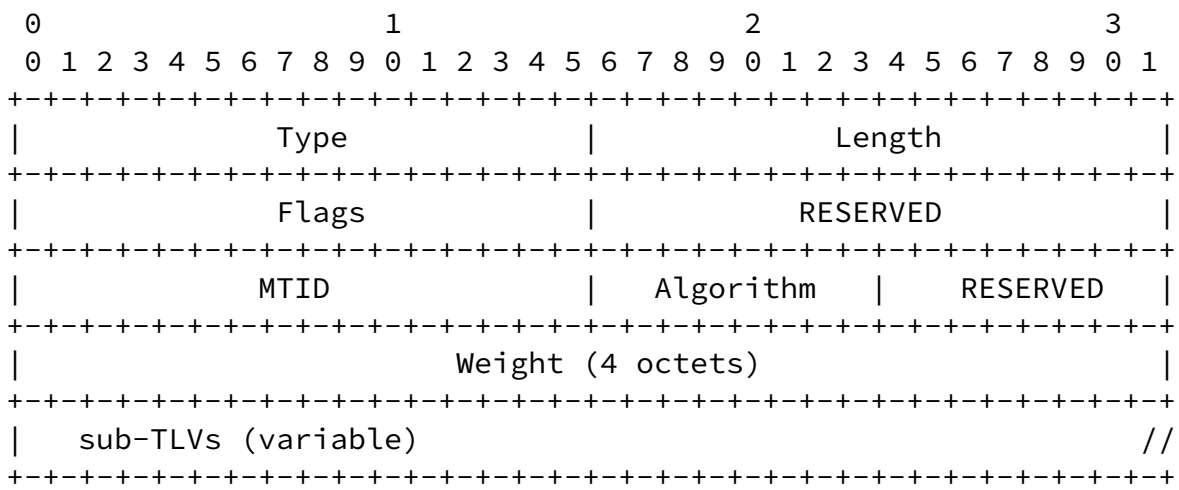

Figure 5. Reverse Segment List Sub-TLV

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

The SR Segment sub-TLV [I-D.ietf-idr-te-lsp-distribution] 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-te-lsp-distribution].

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 [RFC7752] 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, and it can be any applications such as performance measurement [I-D.gandhi-spring-udp-pm] and path re-coputation or re-optimization, etc. The operation of sending information to other precesses is out of scope of this document.

5. IANA Considerations

5.1. 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

6. Security Considerations

TBA

7. 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

8. Acknowledgements

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

9. References

9.1. Normative References

[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", draft-ietf-idr-bgpls-srv6-ext-09 (work in progress), November 2021.

[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-idr-te-lsp-distribution]
Previdi, S., Talaulikar, K., Dong, J., Chen, M., Gredler, H., and J. Tantsura, "Distribution of Traffic Engineering (TE) Policies and State using BGP-LS", draft-ietf-idr-te-lsp-distribution-16 (work in progress), October 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>.

[RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016,

<https://www.rfc-editor.org/info/rfc7752>.

[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>.

9.2. 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.

[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", draft-ietf-idr-sr-policy-path-segment-05 (work in progress), 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

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