IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 17 April 2026

Y. Liu
S. Peng
ZTE
Z. Li
China Mobile

14 October 2025

Supplement of BGP-LS Distribution for SR Policies and State draft-lp-idr-bgp-ls-sr-policy-supplement-03

Abstract

This document supplements additional information of the segment list in the BGP-LS advertisement for SR Policy state information. Two new flags and a new sub-TLV are introduced in the SR Segment List TLV of BGP-LS SR Policy Candidate Path NLRI.

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 17 April 2026.

Copyright Notice

Table of Contents

1. Introduction

SR Policy architecture details are specified in [RFC9256]. An SR Policy comprises one or more candidate paths (CP) of which at a given time one and only one may be active. Each CP in turn may have one or more SID-List of which one or more may be active; when multiple are active then traffic is load balanced over them.

[I-D.ietf-idr-bgp-ls-sr-policy] describes a mechanism to collect the SR policy information that is locally available in a node and advertise it into BGP Link State (BGP-LS) updates. Various TLVs are defined to enable the headend to report the state at the candidate path level and the segment list level.

Currently, a few segment-list-related information is not yet included in [I-D.ietf-idr-bgp-ls-sr-policy]:

* Whether the segment list is a backup path. [I-D.ietf-pce-multipath] proposes extensions to PCEP to specify the protection relationship among segment lists within the candidate path. There would be segment lists in the CP acting as backup for one or more primary segment lists, the backup lists only carry rerouted traffic after the protected path fails.

* Whether the segment list is in administrative shut state. For the candidate path, there's already an S-Flag in the SR Candidate Path State TLV in [I-D.ietf-idr-bgp-ls-sr-policy] indicating the CP is in an administrative shut state. In some usecases, the segment list may also be shut by an administrator for traffic engineering or power saving purpose, e.g, the network administrator may shut certain segment list when the load on the SR Policy is light. This information may also be needed and reported via BGP-LS.

Besides, [RFC8662] describes how Entropy labels (ELs) are to be applied to SR-MPLS to improve load-balancing. Multiple <ELI, EL> pairs may be inserted in the SR-MPLS label stack. A typical use case is that the ingress router (e.g., the headend node) computes a hash based on several fields from a given packet and places the result in the EL(s). The values and the positions of the ELs inserted in the SID-lists may be required by the controller when the headend reports the state of SR Policies via BGP-LS. However, carrying MPLS labels in BGP-LS that are not SR-MPLS SIDs are not yet supported.

This document supplements some additional information of the segment list state as mentioned above in the BGP-LS advertisement for SR Policy state information.

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] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. BGP-LS Extensions for Distributing Segment List States

3.1. New Flags in SR Segment List TLV

SR Segment List TLV is defined in [I-D.ietf-idr-bgp-ls-sr-policy] to report the SID-List(s) of a candidate path. As show in Figure 1,this document introduces two new flags in the flag field of SR Segment List TLV, where,

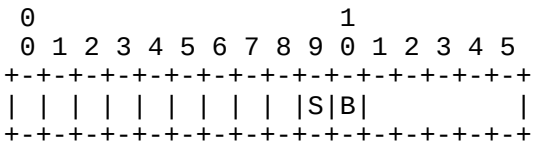

Figure 1: New Flags in the Flag Field of SR Segment List TLV

* S-Flag: Indicates the segment list is in administrative shut state when set. The segment list may be shut by the administrator via CLI or other methods, and it is out of the scope of this document.

* B-Flag: Indicates that the segment list is a pure backup path as specified in [I-D.ietf-pce-multipath] section 4.4 when set. When B-Flag is clear, it indicates it is the primary path that carries normal traffic.

3.2. MPLS Label Sub-TLV

The MPLS Label sub-TLV is defined in this section to carry the generic MPLS Label information. The MPLS Label Sub-TLV is an optional sub-TLV of SR Segment List TLV, and it may be used as the sub-TLV of other TLVs, for the latter case, the detailed usage is out of the scope of this document.

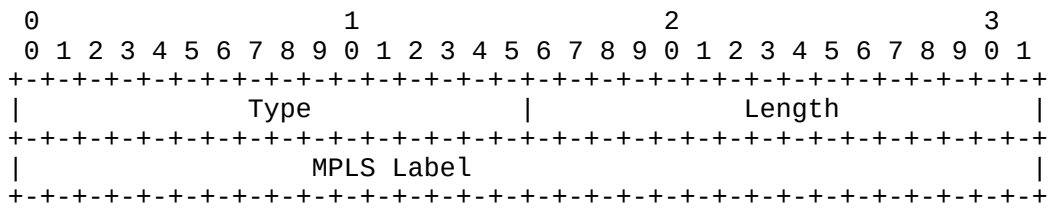

Figure 2: MPLS Label Sub-TLV

Type: TBA

Length: The value is 4 octets.

MPLS Label: a 4-octet-field carrying the MPLS Label.

When receiving the BGP-LS advertisement for the SR Policy Candidate Path NLRI with SR Segment List TLV carrying the MPLS Label sub-TLV, it indicates that there’s a MPLS Label inserted in the SID LIST. The MPLS Label sub-TLV can appear multiple times in the SR Segment List TLV. For example, if there’re two adjacent MPLS Label sub-TLVs in the SR Segment List TLV, with the value of the MPLS Label in the first MPLS Label sub-TLV set as 7 (the reserved label value for ELI[RFC6790]), it indcates that a <ELI, EL> pair has been inserted in the SID list, and the second MPLS Label sub-TLV carries the corresponding EL value.

4. IANA Considerations

This document requests bit 9 and bit 10 in the flag field of "SR Segment List TLV" [I-D.ietf-idr-bgp-ls-sr-policy] under the "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry.

Bit	Description	Reference
9	Administrative Shut State Flag(S-Flag)	This document
10	Backup Path State Flag(B-Flag)	This document

This document requests a new type sub-TLV of "SR Segment List TLV" [I-D.ietf-idr-bgp-ls-sr-policy] under the "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry.

Type	Description	Reference
TBA	MPLS Label 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-bgp-ls-sr-policy].

6. References

6.1. 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>.

[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>.

6.2. Informative References

[I-D.ietf-pce-multipath]
Koldychev, M., Sivabalan, S., Saad, T., Beeram, V. P., Bidgoli, H., Yadav, B., Peng, S., Mishra, G. S., and S. Sidor, "Path Computation Element Communication Protocol (PCEP) Extensions for Signaling Multipath Information", Work in Progress, Internet-Draft, draft-ietf-pce-multipath-14, 5 September 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-pce-multipath-14>.

[RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, DOI 10.17487/RFC6790, November 2012, <https://www.rfc-editor.org/info/rfc6790>.

[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>.

[RFC8662] Kini, S., Kompella, K., Sivabalan, S., Litkowski, S., Shakir, R., and J. Tantsura, "Entropy Label for Source Packet Routing in Networking (SPRING) Tunnels", RFC 8662, DOI 10.17487/RFC8662, December 2019, <https://www.rfc-editor.org/info/rfc8662>.

[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>.

Authors' Addresses

Yao Liu
ZTE
Nanjing
China
Email: liu.yao71@zte.com.cn

Shaofu Peng
ZTE
Nanjing
China
Email: peng.shaofu@zte.com.cn

Zhenqiang Li
China Mobile
Email: lizhenqiang@chinamobile.com