

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 22 January 2026

K. Zhang
J. Dong
Huawei
K. Talaulikar
Cisco Systems
21 July 2025

BGP SR Policy Extensions for metric
draft-ietf-idr-sr-policy-metric-03

Abstract

SR Policy candidate paths can be represented in BGP UPDATE messages. BGP can then be used to propagate the SR Policy candidate paths to the headend nodes in the network. After SR Policy is installed on the ingress node, the packets can be steered into SR Policy through route selection. Therefore, route selection may be performed on the ingress node of the SR Policy. If there are multiple routes to the same destination, the route selection node can select routes based on the local policy. The local policy may use the IGP metric of the selected path, which is the IGP Metric of the SR Policy. Thus the BGP UPDATE message needs to carry the metric of each segment list of the SR Policy Candidate Path, which can be used in path selection of routing.

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 22 January 2026.

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. Motivation	3
4. SR Policy and Tunnel Encapsulation Attribute Update	3
4.1. Metric sub-TLV	4
5. Metric process of SR Policy segment list	5
6. Acknowledgements	6
7. IANA Considerations	6
7.1. New Registry: Metric Sub-TLV	6
7.2. New Registry: Metric Type	6
8. Security Considerations	7
9. References	7
Authors' Addresses	8

1. Introduction

[I-D.ietf-idr-sr-policy-safi] defines SR Policy and Tunnel Encapsulation Attributes. It defines the segment list of the SR policies. Each segment list of an SR Policy is a segment routing path, which may be calculated by path computation element and delivered to the head node of the device by BGP Update Message. On the ingress node, when steering traffic to an SR Policy, the ingress node may need to select between multiple SR Policy paths. And the selection policy may also need the path metric information. Therefore, BGP needs to carry the metric of each path when delivering the segment list of the SR Policy through Update messages to facilitate route selection on the device.

2. Terminology

The following terminology is used in this document.

SR Policy: An ordered list of segments.


Candidate Path: the unit for signaling of an SR Policy to a headend via protocol PCEP or BGP, which is defined in [I-D.ietf-pce-segment-routing-policy-cp] and [I-D.ietf-idr-sr-policy-safi].

SRPM: SR Policy Module.

3. Motivation

In route selection scenarios, the metric of the SR Policy segment list may be required.

The specific scenarios are as follows:



On PE1, the route prefix to CE1 has two different next hop, PE2 and PE3. The next hop to PE1 uses an SR Policy1 on PE1, the endpoint of SR Policy1 is PE2. The next hop to PE2 uses an SR Policy2 on PE1, the endpoint of SR Policy2 is PE3. The prefix to CE1 wants to choose a next hop based on the IGP metric of the route PE1 to PE2 and PE1 and PE3, which uses SR Policy1 and SR Policy2. Thus, BGP needs to pass the IGP metric of SR Policy segment list on PE1.

4. SR Policy and Tunnel Encapsulation Attribute Update

As the metric is defined, the tunnel attribute encapsulation of the BGP SR Policy needs to be updated.

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

Policy Candidate Path Name

Explicit NULL Label Policy (ENLP)

Segment List

Weight

Metric

Segment

Segment

....

....

Where metric indicates the metric for the segment list.

4.1. Metric sub-TLV

A new sub-TLV called Metric sub-TLV is defined. Metric sub-TLV specifies the metric of an SR policy Segment List. Each sub-TLV is encoded as shown in Figure 1. More than one metric Sub-TLVs may be present in one Segment List to refer to the metric values of different metric types.




Figure 1: Metric Sub-TLV

* Type: Metric, 1 octet, TBD.

* Length: 6 octets.

* Metric Type: 1-octet field which identifies the type of the metric being used. The metric-type code points are listed in Section 7.2.

* Flags: None are defined at this stage. Flags SHOULD be set to zero on transmission and MUST be ignored on receipt.

* Metric Value: 4-octet value which indicates the metric of the computed path.

5. Metric process of SR Policy segment list

When the SR Policy headend gets the SR Policy segment list with the metric field, the metric may be of any of the defined types: (IGP metric, Min Unidirection Link delay, TE metric, Hop Count, or SID List length).

The rules for processing SR Policy metrics are as follows:

1. The type of metric to use is determined by the local policy, which can be user-configured. For example, the user specifies the IGP metric as local policy.
2. The metric of the Active Candidate Path is used as the metric of the SR Policy.
3. The metric of the Active Candidate Path uses the maximum metric value of the specified metric type in all segment lists.

Example:

SR Policy: (Headend:1::1, Color: 2, EndPoint: 2::2)
Candidate path preference: 200
Segment list 1: (IGP metric: 20, Link delay: 10, TE metric: 10)
Segment list 2: (IGP metric: 30, Link delay: 20, TE metric: 15)
Candidate Path preference: 100
Segment list 1: (IGP metric: 40, Link delay: 20, TE metric: 20)
Segment list 2: (IGP metric: 30, Link delay: 10, TE metric: 15)
Local policy: IGP metric
Active candidate path: preference 200
Active candidate path metric: 30, which is the maximum IGP metric
of all segment lists in the candidate path.

6. Acknowledgements

TBD.

7. IANA Considerations**7.1. New Registry: Metric Sub-TLV**

This document defines a new Sub-TLV in requests registries "SR Policy List Sub-TLVs" [I-D.ietf-idr-sr-policy-safi]:

Value	Description	Reference
TBD	Metric	This document

Figure 2: Metric sub-TLV

7.2. New Registry: Metric Type

This document requests IANA to maintain a new registry under the "BGP Tunnel Encapsulation" registry. The new registry is called "Metric Type" and contains the codepoints allocated to the "metric type" field defined in Section 4.1. The registry contains the following codepoints, with initial values, to be assigned by IANA with the reference set to this document:

Code Point	Metric Type
0	IGP Metric
1	Min Unidirectional Link Delay [RFC7471]
2	TE Metric [RFC3630]
3	Hop Count (refer [RFC5440])
4	SID List Length
5	Loss Rate
6-250	Unassigned
251-255	Private Use (not to be assigned by IANA)

Figure 3: Metric Type Code Point

8. Security Considerations

The BGP SR policy extension specified in this document is based on [I-D.ietf-idr-sr-policy-safi], the security considerations is the same with align with [I-D.ietf-idr-sr-policy-safi], this draft does not introduce new security requirements.

9. 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-pce-segment-routing-policy-cp]
Koldychev, M., Sivabalan, S., Sidor, S., Barth, C., Peng, S., and H. Bidgoli, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing (SR) Policy Candidate Paths", Work in Progress, Internet-Draft, draft-ietf-pce-segment-routing-policy-cp-27, 4 April 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-segment-routing-policy-cp-27>>.

- [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>>.

Authors' Addresses

Ka Zhang
Huawei
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China
Email: zhangka@huawei.com

Jie Dong
Huawei
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China
Email: jie.dong@huawei.com

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