

IDR
Internet Draft
Intended status: Standards Track
Expires: August 08, 2026

W. Jiang
China Mobile
C. Lin
New H3C Technologies
R. Chen
ZTE Corporation
February 06, 2026

BGP Extensions of SR Policy for Composite Candidate Path
draft-jiang-idr-sr-policy-composite-path-04

Abstract

Segment Routing is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. An SR Policy is associated with one or more candidate paths. A candidate path is either dynamic, explicit or composite. This document defines extensions to BGP to distribute SR policies carrying composite candidate path information. So that composite candidate paths can be installed when the SR policy is applied.

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 August 08, 2026.

Copyright Notice

Copyright (c) 2026 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
1.1. Requirements Language.....	3
2. Constituent SR Policy Attributes in SR Policy.....	3
2.1. Constituent SR Policy Sub-TLV.....	4
2.2. Per-Flow Forwarding Class Sub-TLV.....	5
3. Procedures.....	6
4. Error Handling.....	7
5. Security Considerations.....	7
6. IANA Considerations.....	8
7. References.....	8
7.1. Normative References.....	8
7.2. Informative References.....	9
Authors' Addresses.....	10

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 [RFC9256]. In order to distribute SR policies to the headend, [RFC9830] specifies a mechanism by using BGP.

An SR Policy is associated with one or more candidate paths. A composite candidate path acts as a container for grouping of SR Policies. As described in section 2.2 in [RFC9256], the composite candidate path construct enables combination of SR Policies, for a load-balanced steering of packet flows over its constituent SR Policies.

This document defines extensions to Border Gateway Protocol (BGP) to distribute SR policies carrying composite candidate path information. After BGP distributions valid information about the composite path, the SR Policy Module (SRPM) will instantiate valid SR Policies.

1.1. 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.

2. Constituent SR Policy Attributes in SR Policy

As defined in section 2.2 of [RFC9830], 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
      SRv6 Binding SID
      Preference
      Priority
      Policy Name
      Policy Candidate Path Name
      Explicit NULL Label Policy (ENLP)
      Segment List
        Weight
        Segment
        Segment
        ...
    ...
```

Figure 1: SR Policy Encoding

As described in section 2.2 in [RFC9256], the endpoints of the constituent SR Policies and the parent SR Policy MUST be identical, and the colors of each of the constituent SR Policies and the parent SR Policy MUST be different. Therefore a constituent SR Policy is referenced only by color in the composite candidate path since its headend and endpoint are identical to the parent SR policy.

SR policy with composite candidate 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
 - SRv6 Binding SID
 - Preference
 - Priority
 - Policy Name
 - Policy Candidate Path Name
 - Explicit NULL Label Policy (ENLP)
 - Segment List
 - Weight
 - Segment
 - Segment
 - ...
 - Constituent SR Policy
 - Color
 - Weight
 - Forwarding Class

Figure 2: SR policy with composite candidate path Encoding

SR Policy Architecture [RFC9256] defines the concept of a Composite Candidate Path. A regular SR Policy Candidate Path outputs traffic to a set of Segment Lists, while an SR Policy Composite Candidate Path outputs traffic recursively to a set of SR Policies on the same headend.

2.1. Constituent SR Policy Sub-TLV

The Constituent SR Policy sub-TLV is an optional sub-TLV of BGP Tunnel Encapsulation Attribute, and MAY appear multiple times in the SR Policy encoding. The ordering of Constituent SR Policy sub-TLVs does not matter. The Constituent SR Policy sub-TLV MAY contain a Weight sub-TLV.

The Constituent SR Policy sub-TLV and the Segment List sub-TLV MUST NOT appear in the same candidate path.

The Constituent SR Policy sub-TLV has the following format:

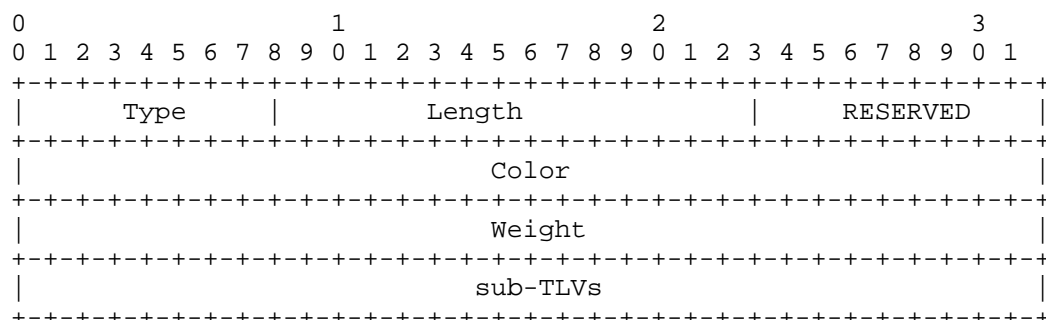


Figure 3: Constituent SR Policy sub-TLV format

where:

- * Type: to be assigned by IANA.
- * Length: The total length (not including the Type and Length fields) of the sub-TLVs encoded within the Constituent SR Policy sub-TLV in terms of octets.
- * RESERVED: 1 octet of reserved bits. This field MUST be set to Zero on transmission and MUST be ignored on receipt.
- * Color: 4 octets that carry an unsigned non-zero integer value indicating the Color of the Constituent SR Policy. As described in section 2.2 in [RFC9256], the endpoints of the constituent SR Policies and the parent SR Policy MUST be identical, thus different constituent SR Policies can be distinguished by Color.
- * Weight: 4 octets carrying an unsigned integer value indicating the weight associated with a segment list as described in Section 2.11 of [RFC9256]. A weight value of zero is invalid.
- * sub-TLVs currently defined: An optional single Per-Flow Forwarding Class sub-TLV which is defined in section 2.2 on this document. The other Sub-TLVs in Constituent SR Policy Sub-TLV are out of scope of this document.

2.2. Per-Flow Forwarding Class Sub-TLV

Per-Flow Forwarding Path builds on top of the concept of the Composite Candidate Path. Each Path in a Per-Flow Forwarding Path is assigned a 3-bit Forward Class (FC) value, which allows QoS classified traffic to be steered depending on the FC. The Per-Flow FC sub-TLV is an optional sub-TLV of the Constituent SR Policy TLV.

The Per-flow FC sub-TLV is OPTIONAL and it MUST NOT appear more than once inside the Constituent SR Policy sub-TLV.

The Per-flow FC sub-TLV has the following format:




Figure 4 Per-Flow FC Sub-TLV

where:

- * Type: to be assigned by IANA.
- * Length: Specifies the length of the value field (not including Type and Length fields) in terms of octets. The value MUST be 6.
- * Flags: 1 octet of flags. No flags are defined in this document. The Flags field MUST be set to zero on transmission and MUST be ignored on receipt.
- * RESERVED-1: 1 octet of reserved bits. This field MUST be set to zero on transmission and MUST be ignored on receipt.
- * Reserved-2(29 bits): This field MUST be set to zero on transmission and MUST be ignored on receipt.
- * FC (3 bits): Forward class value that is given by the QoS classifier to traffic entering the given Candidate Path. Different classes of traffic that enter the given Candidate Path can be differentially steered into different Colors.

3. Procedures

The document does not bring new operation beyond the description of operations defined in [RFC9830]. The existing operations defined in [RFC9830] can apply to this document directly.

Typically but not limit to, the SR policies carrying composite candidate path information are configured by a controller.

After configuration, the SR policies carrying path composite candidate path information will be advertised by BGP update messages.

The operation of advertisement is the same as defined in [RFC9830], as well as the reception.

Note the differences among several color TLVs. The Tunnel Egress Endpoint and Color sub-TLVs of the Tunnel Encapsulation Attribute, as defined in [RFC9012], are not utilized for SR Policy encodings; see more details in Section 2.3 of [RFC9830].

The Color Extended Community (as defined in [RFC9012]) is used to steer traffic into an SR Policy, as described in Section 8.8 of [RFC9256] and Section 3 of [RFC9830].

The color of the Constituent SR Policy is identified by its color, as described in Section 2.1.

4. Error Handling

The error handling of the BGP Update messages for BGP SR Policy SAFI with the NRP extensions defined in this document follows the procedures in section 5 of [RFC9830].

The validation of the TLVs/sub-TLVs introduced in this document and defined in their respective subsections of Section 2 MUST be performed to determine if they are malformed or invalid.

The Constituent SR Policy sub-TLV and the Segment List sub-TLV MUST NOT appear in the same candidate path. If Constituent SR Policy sub-TLV does not match the above description, or its format is considered malformed, the associated BGP SR Policy NLRI is considered malformed and the "treat-as-withdraw" strategy of [RFC7606] MUST be applied.

The Per-flow FC sub-TLV is optional and MUST NOT appear more than once for one Constituent SR Policy sub-TLV. The Per-flow FC sub-TLV is considered malformed if its format does not match the above description. If the Per-flow FC sub-TLV appears more than once, or its format is considered malformed, the associated BGP SR Policy NLRI is considered malformed and the "treat-as-withdraw" strategy of [RFC7606] MUST be applied.

5. Security Considerations

The security considerations of BGP [RFC4271] and BGP SR policy [RFC9830] apply to this document.

This document defines BGP extensions for distributing SR policies that carry composite candidate path information. These functions extend the risks associated with SR Policy into the dynamic realm. Misconfiguration or errors in configuring an SR Policy Composite Candidate Path may lead to packets being forwarded along unintended paths for the affected routes.

6. IANA Considerations

This document defines a sub-TLV in the "BGP Tunnel Encapsulation Attribute Sub-TLVs" registry under the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group.

Value	Description	Reference
TBA	Constituent SR Policy Sub-TLV	This document

This document creates a new registry called "Constituent SR Policy sub-TLV" under the "Border Gateway Protocol (BGP) Tunnel Encapsulation" registry group.

Value	Description	Reference
TBA	Per-Flow FC Sub-TLV	This document

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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>>.

- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017
- [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>>.
- [RFC9830] Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", RFC 9830, DOI 10.17487/RFC9830, September 2025, <<https://www.rfc-editor.org/info/rfc9830>>.

7.2. Informative References

- [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

Wenying Jiang
China Mobile
Beijing
China
Email: jiangwenying@chinamobile.com

Changwang Lin
New H3C Technologies
Beijing
China
Email: linchangwang.04414@h3c.com

Ran Chen
ZTE Corporation
Email: chen.ran@zte.com.cn

