

IDR Working Group
Internet Draft
Intended status: Standards Track
Expires: October 27, 2026

C. Lin
New H3C Technologies
Y. Liu
China Mobile
R. Chen
ZTE Corporation
X. Yi
China Unicom
April 25, 2026

BGP Extensions of SR Policy for Headend Behavior
draft-lin-idr-sr-policy-headend-behavior-06

Abstract

[RFC8986] defines H. Encaps behavior, H. Encaps.Red behavior, H. Encaps.L2 behavior, and H. Encaps.L2.Red behavior for SR policy. This document defines extensions to Border Gateway Protocol (BGP) to distribute SR policies carrying headend behavior.

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 October 27, 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. Use Case.....	3
3. Headend Behavior in SR Policy.....	4
3.1. Headend Behavior Sub-TLV.....	5
3.2. L2 Headend Behavior Sub-TLV.....	6
4. Security Considerations.....	7
5. IANA Considerations.....	7
6. References.....	7
6.1. Normative References.....	7
6.2. Informative References.....	8
Authors' Addresses.....	8

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.

As described in [RFC9256], a headend can steer a packet flow into an SR Policy in various ways, including BSID steering, per-destination steering, per-flow steering, and policy-based steering. Moreover, [I-D.ietf-idr-ts-flowspec-srv6-policy] describes a way by using BGP FlowSpec to steer packets into an SRv6 Policy.

[RFC8986] defines End.B6.Encaps behavior and End.B6.Encaps.Red behavior for SRv6 BSID. When receiving packets with an active SID matching a local BSID of these kinds, the headend will perform corresponding behaviors. Different BSID behaviors are suitable for different scenarios. For example, comparing with End.B6.Encaps, End.B6.Encaps.Red reduces the size of the SRH by excluding the first SID, which can be useful for the devices with lower capacity of SID depths, like the switches in data center network.

The SRv6 Binding SID sub-TLV is defined in [RFC9830] to signal the SRv6 BSID information along with SR Policies. It enables the specified SRv6 BSID behavior to be instantiated on the headend node. However, if the packets are steering into an SR Policy in some other way than using BSID, the headend behavior is not specified during the distributing of SR Policy by BGP. The network operator has to use additional tools, like NETCONF, to signal the headend behavior.

[RFC8986] defines H. Encaps behavior, H. Encaps.Red behavior, H. Encaps.L2 behavior, and H. Encaps.L2.Red behavior for SR policy. This document defines extensions to Border Gateway Protocol (BGP) to distribute SR policies carrying headend behavior. So that the headend can be instructed to perform specific behavior when packets are steered into the SR policy without BSID.

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. Use Case

[RFC8986] defines describes a set of SRv6 Policy Headend [RFC8402] behaviors.

- o H.Encaps: SR Headend with Encapsulation in an SR Policy
- o H.Encaps.Red: H.Encaps with Reduced Encapsulation
- o H.Encaps.L2: H.Encaps Applied to Received L2 Frames
- o H.Encaps.L2.Red: H.Encaps.Red Applied to Received L2 Frames

When a source does not need to preserve the entire SID list within the Segment Routing Header (SRH), a reduced SRH may be utilized. This reduced SRH (see Section 4.1.1 of [RFC8754]) omits the first segment of the associated SR Policy, as it already appears in the Destination Address (DA) of the IPv6 header. Utilizing the NEXT-CSID method from [RFC9800] enables further reduction in SRH size, thereby improving compression efficiency. In cases where traceability is necessary, a complete SID list should be preserved, which can be achieved through the Encap mode. Additionally, by leveraging the NEXT-CSID and/or REPLACE-CSID flavors as specified in [RFC9800] to minimize packet header length, the reduced SRH can further enhance compression performance.

3. Headend Behavior in SR Policy

As defined in [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

 ...

 ...

SR policy with headend behavior is expressed 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)

Headend Behavior

L2 Headend Behavior

Segment List

Weight

Segment

Segment

...

...

3.1. Headend Behavior Sub-TLV

The Headend Behavior sub-TLV encodes the default headend behavior associated with the candidate path for L3 traffic. When the headend steers L3 packets into that SR Policy and the associated candidate path is active, the specific headend behavior should be performed by default. In the case of BSID steering, the behavior defined by the BSID overrides the default headend behavior.

The Headend Behavior sub-TLV is optional, and MUST NOT appear more than once in the SR Policy encoding.

The Headend Behavior sub-TLV has the following format:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
Type										Length										RESERVED																			
Headend Behavior																																							

where:

- o Type: to be assigned by IANA.
- o Length: 4.

- o RESERVED: 2 octets of reserved bits. SHOULD be set to zero on transmission and MUST be ignored on receipt.
- o Headend Behavior: a 2-octet value. The following values are defined.
 - * Headend Behavior = 0: H.Encaps. A headend behavior defined in [RFC8986].
 - * Headend Behavior = 1: H.Encaps.Red. A headend behavior defined in [RFC8986].

3.2. L2 Headend Behavior Sub-TLV

The L2 Headend Behavior sub-TLV encodes the default headend behavior associated with the candidate path for L2 traffic. When the headend steers L2 packets into that SR Policy and the associated candidate path is active, the specific headend behavior should be performed by default.

The L2 Headend Behavior sub-TLV is optional, and MUST NOT appear more than once in the SR Policy encoding.

The L2 Headend Behavior sub-TLV has the following format:

```

      0                   1                   2                   3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
      +-----+-----+-----+-----+-----+-----+-----+-----+
      |          Type          |      Length      |      RESERVED      |
      +-----+-----+-----+-----+-----+-----+-----+-----+
      |      L2 Headend Behavior      |
      +-----+-----+-----+-----+-----+-----+-----+
  
```

where:

- o Type: to be assigned by IANA.
- o Length: 4.
- o RESERVED: 2 octets of reserved bits. SHOULD be set to zero on transmission and MUST be ignored on receipt.
- o L2 Headend Behavior: a 2-octet value. The following values are defined.
 - * L2 Headend Behavior = 0: H.Encaps.L2. A headend behavior defined in [RFC8986].

- * L2 Headend Behavior = 1: H.Encaps.L2.Red. A headend behavior defined in [RFC8986].

4. Security Considerations

Procedures and protocol extensions defined in this document do not affect the security considerations discussed in [RFC9830].

5. IANA Considerations

Headend Behavior Sub-TLV (TBD)

L2 Headend Behavior Sub-TLV (TBD)

6. References

6.1. Normative References

- [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>>.
- [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>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header(SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [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>>.

[RFC9800] Cheng, W., Ed., Filsfils, C., Li, Z., Decraene, B., and F. Clad, Ed., "Compressed SRv6 Segment List Encoding", RFC 9800, DOI 10.17487/RFC9800, June 2025, <<https://www.rfc-editor.org/info/rfc9800>>.

[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>>.

6.2. Informative References

[I-D.ietf-idr-ts-flowspec-srv6-policy] Jiang, W., Liu, Y., Zhuang, S., Mishra, G., and S. Chen, "Traffic Steering using BGP Flowspec with SR Policy", Work in Progress, Internet-Draft, draft-ietf-idr-ts-flowspec-srv6-policy-07, 04 August 2025, <<http://www.ietf.org/internet-drafts/draft-ietf-idr-ts-flowspec-srv6-policy-07.txt>>.

Authors' Addresses

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

Yisong Liu
China Mobile
China
Email: liuyisong@chinamobile.com

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

Xinxin Yi
China Unicom
Beijing, 100048
China
Email: yixx3@chinaunicom.cn

