

IDR Working Group  
Internet Draft  
Intended status: Standards Track  
Expires: June 5, 2024

C. Lin  
New H3C Technologies  
W. Jiang  
Y. Liu  
China Mobile  
M. Chen  
H. Li  
New H3C Technologies  
December 8, 2023

BGP Extensions of SR Policy for Headend Behavior  
draft-lin-idr-sr-policy-headend-behavior-03

Abstract

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 June 5, 2024.

Copyright Notice

Copyright (c) 2023 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 (<http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

1. Introduction.....	2
1.1. Requirements Language.....	3
2. Headend Behavior in SR Policy.....	3
2.1. Headend Behavior Sub-TLV.....	4
2.2. L2 Headend Behavior Sub-TLV.....	5
3. Extensions of BGP-LS.....	6
4. Security Considerations.....	6
5. IANA Considerations.....	6
6. References.....	7
6.1. Normative References.....	7
6.2. Informative References.....	7
Authors' Addresses.....	9

## 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, [I-D.ietf-idr-segment-routing-te-policy] 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.jiang-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. [I-D.filsfils-spring-srv6-net-pgm-insertion] extends the SRv6 BSID behaviors with End.B6.Insert and End.B6.Insert.Red. 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. End.B6.Insert inserts a new SRH in between the IPv6 Header and the received SRH rather than pushing a new IPv6 header, which can be applied to

express scalable traffic-engineering policies across multiple domains.

The SRv6 Binding SID sub-TLV is defined in [I-D.ietf-idr-segment-routing-te-policy] 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.

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. Headend Behavior in SR Policy

As defined in [I-D.ietf-idr-segment-routing-te-policy], 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

...

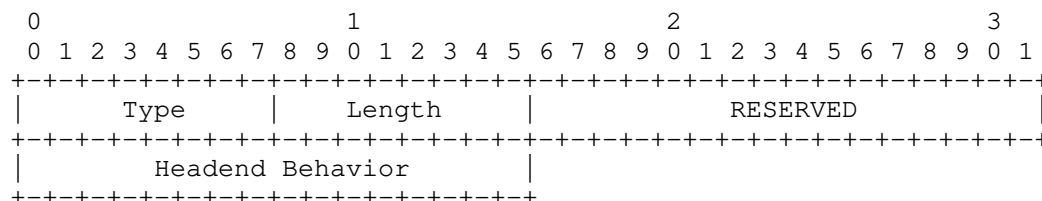
...

## 2.1. 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:



where:

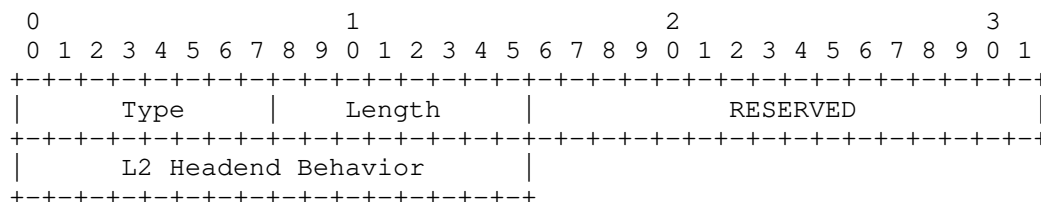
- o Type: to be assigned by IANA.
- o Length: 4.
- o RESERVED: 2 octets of reserved bits. SHOULD be set to zero on transmission and MUST be ignored on receipt.
- o Headend Behavior: a 2-octet value. The following values are defined.
  - \* TBD: H.Encaps. A headend behavior defined in [RFC8986].
  - \* TBD: H.Encaps.Red. A headend behavior defined in [RFC8986].
  - \* TBD: H.Insert. A headend behavior defined in [I-D.filsfils-spring-srv6-net-pgm-insertion].
  - \* TBD: H.Insert.Red. A headend behavior defined in [I-D.filsfils-spring-srv6-net-pgm-insertion].

## 2.2. 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:



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.
  - \* TBD: H.Encaps.L2. A headend behavior defined in [RFC8986].
  - \* TBD: H.Encaps.L2.Red. A headend behavior defined in [RFC8986].

### 3. Extensions of BGP-LS

[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-LS updates. Extensions of BGP-LS for headend behavior of SR Policy will be included in the future version of this draft.

### 4. Security Considerations

Procedures and protocol extensions defined in this document do not affect the security considerations discussed in [I-D.ietf-idr-segment-routing-te-policy].

### 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] Filts, 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>>.
- [RFC9256] Filts, 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>>.
- [I-D.ietf-idr-segment-routing-te-policy] Previdi, S., Filts, C., Talaulikar, K., Mattes, P., Rosen, E., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-segment-routing-te-policy-26, 23 October 2023, <<http://www.ietf.org/internet-drafts/draft-ietf-idr-segment-routing-te-policy-26.txt>>.

### 6.2. Informative References

- [RFC8986] Filts, 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>>.
- [I-D.jiang-idr-ts-flowspec-srv6-policy] Jiang, W., Liu, Y., Chen, S., and S. Zhuang, "Traffic Steering using BGP Flowspec with SRv6 Policy", Work in Progress, Internet-Draft, draft-jiang-idr-ts-flowspec-srv6-policy-07, 23 March 2022, <<http://www.ietf.org/internet-drafts/draft-jiang-idr-ts-flowspec-srv6-policy-07.txt>>.

[I-D.filsfils-spring-srv6-net-pgm-insertion] Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "SRv6 NET-PGM extension: Insertion", Work in Progress, Internet-Draft, draft-filsfils-spring-srv6-net-pgm-insertion-09, 16 August 2023, <<http://www.ietf.org/internet-drafts/draft-filsfils-spring-srv6-net-pgm-insertion-09.txt>>.

[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-03, 5 November 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-03>>.

Authors' Addresses

Changwang Lin  
New H3C Technologies

Email: [linchangwang.04414@h3c.com](mailto:linchangwang.04414@h3c.com)

Wenying Jiang  
China Mobile

Email: [jiangwenying@chinamobile.com](mailto:jiangwenying@chinamobile.com)

Yisong Liu  
China Mobile

Email: [liuyisong@chinamobile.com](mailto:liuyisong@chinamobile.com)

Mengxiao Chen  
New H3C Technologies

Email: [chen.mengxiao@h3c.com](mailto:chen.mengxiao@h3c.com)

Hao Li  
New H3C Technologies

Email: [lihao@h3c.com](mailto:lihao@h3c.com)

