

IDR
Internet Draft
Intended status: Standards Track
Expires: August 18, 2025

C. Lin
New H3C Technologies
J. Li
China Mobile
Ran.Chen
ZTE
February 18, 2025

BGP SR Policy Extensions for Administrative Flags
draft-lin-idr-sr-policy-admin-flags-01

Abstract

Segment Routing is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. An SR Policy is a set of candidate paths, each consisting of one or more segment lists.

This document defines an extension to the BGP SR Policy that sets the administrative state of the candidate path or segment list, facilitating the operation and maintenance of the SR Policy.

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

This Internet-Draft will expire on August 18, 2025.

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 (<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. Admin State in SR Policy.....	3
2.1. Candidate Path Administrative Flags Sub-TLV.....	5
2.2. Segment List Administrative Flags Sub-TLV.....	6
3. Security Considerations.....	6
4. IANA Considerations.....	7
5. References.....	7
5.1. Normative References.....	7
5.2. Informative References.....	7
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, [I-D.ietf-idr-sr-policy-safi] specifies a mechanism by using BGP.

For management purposes, the controller may occasionally need to temporarily divert traffic from a specific forwarding path and then restore it later. In such cases, the controller can issue an Administrative Down command to a specific path in the SR Policy on the device without removing the path. When it is time to restore the path, the controller can simply issue an Administrative Up command to that same path.

In another scenario, such as in 6PE or EPE situations where it is necessary to conserve service route SIDs, the SR Policy Flag can be extended to indicate settings. For example, configuring the Candidate Path as "Ignore service routes Prefix SID" can help optimize the segment list.

In some scenarios, the Candidate Path cannot be used as a backup path. The operator needs to control the Candidate Path status to identify whether this CP can serve as a backup path.

In [RFC9256], section 8.2 defines the Drop-Upon-Invalid behavior. An SR Policy MAY be enabled for the Drop-Upon-Invalid behavior. Currently, there is no behavior control for Drop-Upon-Invalid on the path of an SR policy.

This document defines an extension to the BGP SR Policy that sets the management state of the candidate path or the segment list, facilitating the operation and maintenance of the SR Policy.

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. Admin State in SR Policy

As defined in [I-D.ietf-idr-sr-policy-safi], 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 Administrative Flags are 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

Policy Candidate Path Administrative Flags

Explicit NULL Label Policy (ENLP)

Segment List

Weight

Segment List Administrative Flags

Segment

Segment

...

...

The Candidate Path Administrative Flags can also be advertised using the Candidate Path Administrative Flags sub-TLV, as defined in Section 2.1.

The segment list Administrative Flags can be advertised using the Segment List Administrative Flags sub-TLV, as defined in Section 2.2.

2.1. Candidate Path Administrative Flags Sub-TLV

The Candidate Path Administrative Flags sub-TLV is used to indicate the AdminState of the Candidate Path.

The Candidate Path Administrative Flags sub-TLV is optional and it MUST NOT appear more than once inside the Segment List sub-TLV.

The Candidate Path Administrative Flags 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 2
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type      |      Length      |      Flags      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

- o Type: TBD.
- o Length: 2.
- o Flags: 2 octet of flags.

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|S|B|I|D|      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 5: Administrative Flags

where:

- S-Flag: Indicates the CP is in an administrative shut state when set.
- B-Flag: Indicates the CP is configured as "backup ineligible".
- I-Flag: Indicates the CP is configured as "Ignore service route's Prefix SID". It allows traffic to a BGP service route to be steered over an SR policy without imposing the service route's prefix label or SRv6 Service SID.
- D-Flag: Indicates the CP has been marked as the "Drop Upon Invalid" behavior as described in section 8.2 of [RFC9256].

2.2. Segment List Administrative Flags Sub-TLV

The Segment List Administrative Flags sub-TLV is used to indicate the AdminState of the Segment List of Candidate Path.

The Segment List Administrative Flags sub-TLV is optional and it MUST NOT appear more than once inside the Segment List sub-TLV.

The Segment List Administrative Flags 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 2
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Type      |      Length      |      Flags      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

- o Type: TBD.
- o Length: 2.
- o Flags: 2 octet of flags.

```

0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|S|                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 6: Administrative Flags

where:

- S-Flag: Indicates the CP is in an administrative shut state when set.

3. Security Considerations

The security requirements and mechanisms described in [I-D.ietf-idr-sr-policy-safi] also apply to this document.

This document does not introduce any new security consideration.

4. IANA Considerations

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

Value	Description	Reference
TBD1	Candidate Path Administrative Flags sub-TLV	This document
TBD2	Segment List Administrative Flags sub-TLV	This document

5. References

5.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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>>.
- [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-06, 30 July 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-06>>.

5.2. 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

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

Jinming Li
China Mobile
32 Xuanwumen West Street
Beijing
Xicheng District, 100053
China
Email: lijinming@chinamobile.com

Ran Chen
ZTE Corporation

Email: chen.ran@zte.com.cn

