

Workgroup: IDR Working Group
Internet-Draft:
draft-lp-idr-sr-path-protection-06
Published: 28 June 2023
Intended Status: Standards Track
Expires: 30 December 2023
Authors: Y. Liu S. Peng C. Lin
 ZTE ZTE New H3C Technologies
 M. Koldychev
 Cisco Systems, Inc.
BGP Extensions of SR Policy for Segment List Protection

Abstract

This document proposes extensions of BGP in order to provide protection information for segment lists when delivering SR policy via BGP.

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 30 December 2023.

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 (<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](#)
 - [1.1. Requirements Language](#)
- [2. BGP Extensions for Advertising Segment List](#)
 - [2.1. Extensions of Segment List sub-TLV](#)
 - [2.2. List Protection Sub-TLV](#)
- [3. IANA Considerations](#)
 - [3.1. New Registry: Flag Field of Segment List sub-TLV](#)
 - [3.2. Existing Registry: BGP Tunnel Encapsulation Attribute sub-TLVs](#)
- [4. Security Considerations](#)
- [5. References](#)
 - [5.1. Normative References](#)
 - [5.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

Segment Routing [[RFC8402](#)] allows a headend node to steer a packet flow along any path. [[RFC9256](#)] details the concept of SR Policy and steering into an SR Policy. An SR Policy is a set of candidate paths, each consisting of one or more segment lists. The headend of an SR Policy may learn multiple candidate paths for an SR Policy.

Candidate path can be used for path protection, that is, the lower preference candidate path may be designated as the backup for a specific or all (active) candidate path(s). Backup candidate path provide protection only when all the segment lists in the active CP are invalid. If a candidate path is associated with a set of Segment-Lists, each Segment-List is associated with weight for weighted load balancing.

The protection mechanism for SR Policy is not flexible enough. For example, there're two active segment lists(SL1, SL2) in the primary candidate path CP1, SL1 and SL2 can together carry 80 Gbps. If SL1 fails, CP1 are still the primary path, but the bandwidth of CP1 is probably not enough. If there's a backup segment list for SL1, e.g, SL3, in CP1, traffic will be load-balanced between SL3 and SL2 after SL1 fails.

The pcep extensions for segment list identification and protection relationship among segment lists specification are proposed in [[I-D.ietf-pce-multipath](#)].

[[I-D.ietf-idr-segment-routing-te-policy](#)] specifies BGP extensions for the advertisement of SR Policy.

[[I-D.lin-idr-sr-policy-seglist-id](#)] defines extensions to BGP SR Policy to specify the identifier of segment list.

This document proposes extensions of BGP in order to provide the protection information of segment lists when delivering SR policy via BGP.

1.1. 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 [\[RFC2119\]](#) [\[RFC8174\]](#) when, and only when, they appear in all capitals, as shown here.

2. BGP Extensions for Advertising Segment List

2.1. Extensions of Segment List sub-TLV

Segment List sub-TLV is introduced in [\[I-D.ietf-idr-segment-routing-te-policy\]](#) and it includes the elements of the paths (i.e., segments).

This document introduces a one-bit flag in the RESERVED field, where,




Figure 1: B-Flag in Segment List sub-TLV

*B-Flag(Backup Flag): One bit. If set, indicates a pure backup path. This is a segment list that only carries rerouted traffic after the protected segment list fails. If this flag is not set, it indicates that the segment list acts as the active member in the candidate path that carries normal traffic.

Using segment lists for path protection can be compatible with using candidate paths. When a path fails, the backup segment list within the same candidate path is used preferentially for path protection. If the backup list is also invalid, then other candidate path can be enabled for protection.

2.2. List Protection Sub-TLV

This document introduces a new sub-sub-tlv of Segment List sub-TLV, where,



- *Length: 1 octet, specifies the length of the value field not including Type and Length fields.
- *RESERVED: 2 octet of reserved bits. SHOULD be unset on transmission and MUST be ignored on receipt.
- *Backup List ID: 4 octet of ID for the back up segment list, the segment list id is delivered in Segment List ID Sub-TLV as defined in [[I-D.lin-idr-sr-policy-seglist-id](#)]. If there're multiple backup paths, the list ID of each path should be included in the TLV.

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

Preference

Priority

Policy Name

Explicit NULL Label Policy (ENLP)

Segment List

Weight

Segment

Segment

...

Segment List

...

...

The new SR Policy encoding structure with List Protection sub-TLV is shown 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

List Protection

Weight

Segment

Segment

...

Segment List

...

...

3. IANA Considerations

3.1. New Registry: Flag Field of Segment List sub-TLV

This document introduces a one-bit flag field in the Segment List sub-TLV [[I-D.ietf-idr-segment-routing-te-policy](#)] for the Backup Flag (B-Flag).

3.2. Existing Registry: BGP Tunnel Encapsulation Attribute sub-TLVs

This document defines a new sub-TLV in the registry "SR Policy List Sub-TLVs" [[I-D.ietf-idr-segment-routing-te-policy](#)] to be assigned by IANA:

| Codepoint | Description | Reference |
|-----------|-------------------------|---------------|
| ----- | ----- | ----- |
| TBD | List Protection Sub-TLV | This document |

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. References

5.1. Normative References

[I-D.ietf-idr-segment-routing-te-policy]

Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-segment-routing-te-policy-20, 27 July 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-segment-routing-te-policy-20>>.

[I-D.lin-idr-sr-policy-seglist-id]

Lin, C., Cheng, W., Yao, L., Talaulikar, K., and M. Chen, "BGP SR Policy Extensions for Segment List Identifier", Work in Progress, Internet-Draft, draft-lin-idr-sr-policy-seglist-id-03, 3 April 2023, <<https://datatracker.ietf.org/doc/html/draft-lin-idr-sr-policy-seglist-id-03>>.

[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>>.

5.2. Informative References

[I-D.ietf-pce-multipath]

Koldychev, M., Sivabalan, S., Saad, T., Beeram, V. P., Bidgoli, H., Yadav, B., Peng, S., and G. S. Mishra, "PCEP Extensions for Signaling Multipath Information", Work in Progress, Internet-Draft, draft-ietf-pce-multipath-08, 1 May 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-multipath-08>>.

[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>>.

[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

Yao Liu
ZTE

Email: liu.yao71@zte.com.cn

Shaofu Peng
ZTE

Email: peng.shaofu@zte.com.cn

Changwang Lin
New H3C Technologies

Email: linchangwang.04414@h3c.com

Mike Koldychev
Cisco Systems, Inc.

Email: mkoldych@cisco.com