

IDR WG
Internet-Draft
Intended status: Standards Track
Expires: 22 June 2022

Y. Liu
S. Peng
ZTE
19 December 2021

BGP Extensions of SR Policy for Path Protection
draft-lp-idr-sr-path-protection-02

Abstract

This document proposes extensions of BGP to provide protection information of segment lists within a candidate path when delivering SR policy. And it also extends BGP-LS to provide some extra information of the segment list in the advertisement.

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 22 June 2022.

Copyright Notice

Copyright (c) 2021 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.	BGP Extensions for Advertising Segment List	3
2.1.	Extensions of Segment List sub-TLV	3
2.2.	List Identifier Sub-TLV	4
2.2.1.	List Protection Sub-TLV	4
3.	BGP-LS Extensions for Distributing Segment List States	7
4.	IANA Considerations	7
4.1.	New Registry: Flag Field of Segment List sub-TLV	7
4.2.	Existing Registry: BGP Tunnel Encapsulation Attribute sub-TLVs	7
4.3.	New Registry: List Identifier Sub-TLVs	8
4.4.	Existing Registry: Flag Field of SR Segment List TLV	8
5.	Security Considerations	8
6.	References	8
6.1.	Normative References	8
6.2.	Informative References	9
	Authors' Addresses	9

[1.](#) Introduction

Segment Routing [[RFC8402](#)] allows a headend node to steer a packet flow along any path. [[I-D.ietf-spring-segment-routing-policy](#)] 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 three segment lists(SL1, SL2, SL3) in candidate path 1, it may be desired that SL1 and SL2 are the primary path, SL3 are the backup path for SL1 and will be active only when SL1 fails.

[I-D.ietf-pce-multipath] proposes extensions to PCEP to specify the protection relationship between segment lists in the candidate path.

Internet-Draft

BGP Extensions for Segment List

December 2021

[I-D.ietf-idr-segment-routing-te-policy] specifies BGP extensions for the advertisement of SR Policies and each candidate path is carried in an NLRI. This document proposes extensions of BGP in order to provide protection information of segment lists when delivering SR policy.

[I-D.ietf-idr-te-lsp-distribution] describes a mechanism to collect the SR policy information that is locally available in a node and advertise it into BGP Link State (BGP-LS) updates. This document also extends it to provide some extra information of the segment list in a candidate path in the BGP-LS advertisement.

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](#)].

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.

[illegible]

Figure 1: Segment List sub-TLV

B-Flag(Backup Flag): one bit. When set to 0, it indicates that the segment list acts as the active member in the candidate path. When set to 1, it indicates that the segment list acts as the backup path in the candidate path.

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 Identifier Sub-TLV

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

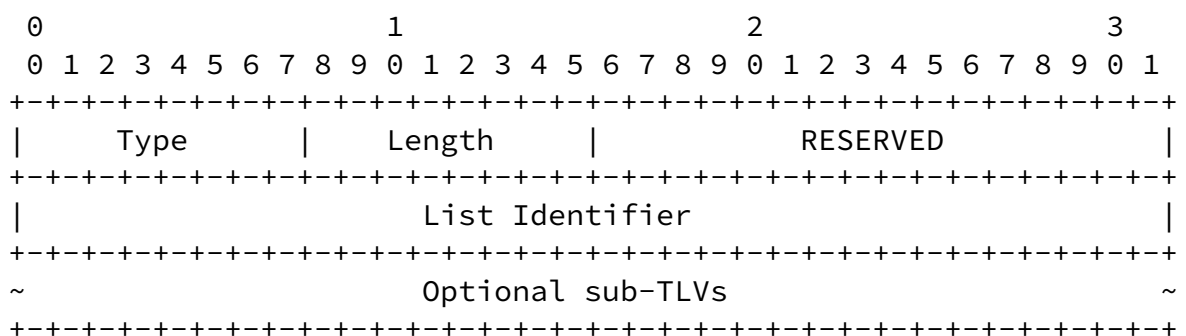


Figure 2: List Identifier Sub-TLV

- * Type: 1 octet. TBD.
- * 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.
- * List Identifier: 4 octets. It is the identifier of the corresponding segment list, so that the segment list can be operated according to the specified Segment List identifier.
- * This sub-TLV is optional and it MUST NOT appear more than once inside the Segment List sub-TLV.

2.2.1. List Protection Sub-TLV

The List Protection Info sub-TLV is an optional sub-TLV of List Identifier sub-TLV, where:

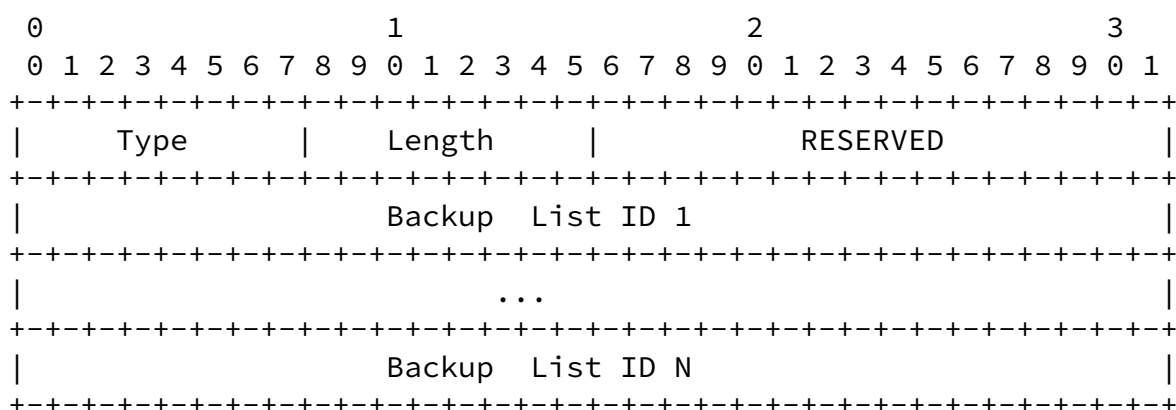


Figure 3: List Protection Info Sub-TLV

- * Type: 1 octet. TBD.
- * 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 octets. It is the List Identifier of the backup segment list that protects this segment list. 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 Identifier 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 Identifier
      List Protection Info
      Weight
      Segment
      Segment
      ...
    Segment List
      ...
    ...
```

[3.](#) BGP-LS Extensions for Distributing Segment List States

[I-D.ietf-idr-te-lsp-distribution] describes a mechanism to collect the SR Policy information that is locally available in a node and advertise it into BGP Link State (BGP-LS) updates. The SR Policy information includes status of the candidate path, e.g, whether the candidate path is administrative shut or not.

SR Segment List TLV is defined in [[I-D.ietf-idr-te-lsp-distribution](#)] to report the SID-List(s) of a candidate path. Figure 4 shows the flags in SR Segment List TLV.




Figure 4: Flag Field of SR Segment List TLV

The D,E,C,V,R,F,A,M flags are defined in [[I-D.ietf-idr-te-lsp-distribution](#)].

This document introduces two new flags, where,

- * S-Flag : Indicates the segment list is in administrative shut state when set.
- * B-Flag : Indicates the segment list is the backup path within the candidate path when set, otherwise it is the active path.

4. IANA Considerations

4.1. 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).

4.2. 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 Identifier Sub-TLV	This document

4.3. New Registry: List Identifier Sub-TLVs

This document requests the creation of a new registry called "List Identifier Sub-TLVs" under the "BGP Tunnel Encapsulation" registry. Following initial Sub-TLV codepoint are assigned by this document.

Codepoint	Description	Reference
TBD	List Protection Sub-TLV	This document

[4.4.](#) Existing Registry: Flag Field of SR Segment List TLV

This document requests bit 9 and bit 10 in the flag field of "SR Segment List TLV" [[I-D.ietf-idr-te-lsp-distribution](#)] under the "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" registry.

Bit	Description	Reference
9	Administrative Shut State Flag(S-Flag)	This document
10	Backup Path State Flag(B-Flag)	This document

[5.](#) 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](#)] and [[I-D.ietf-idr-te-lsp-distribution](#)].

[6.](#) References

[6.1.](#) Normative References

[I-D.ietf-idr-segment-routing-te-policy]
Previdi, S., Filisfilis, 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-14](#), 10 November 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-segment-routing-te-policy-14>>.

[I-D.ietf-idr-te-lsp-distribution]
Previdi, S., Talaulikar, K., Dong, J., Chen, M., Gredler, H., and J. Tantsura, "Distribution of Traffic Engineering (TE) Policies and State using BGP-LS", Work in Progress, Internet-Draft, [draft-ietf-idr-te-lsp-distribution-16](#), 22 October 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-te-lsp-distribution-16>>.

[I-D.ietf-spring-segment-routing-policy]

Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", Work in Progress, Internet-Draft, [draft-ietf-spring-segment-routing-policy-14](https://datatracker.ietf.org/doc/html/draft-ietf-spring-segment-routing-policy-14), 25 October 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-segment-routing-policy-14>>.

[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>>.

[6.2](#). Informative References

[I-D.ietf-pce-multipath]

Koldychev, M., Sivabalan, S., Saad, T., Beeram, V. P., Bidgoli, H., Yadav, B., and S. Peng, "PCEP Extensions for Signaling Multipath Information", Work in Progress, Internet-Draft, [draft-ietf-pce-multipath-03](https://datatracker.ietf.org/doc/html/draft-ietf-pce-multipath-03), 25 October 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-multipath-03>>.

[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>>.

Authors' Addresses

Yao Liu
ZTE
Nanjing
China

Email: liu.yao71@zte.com.cn

Shaofu Peng
ZTE
Nanjing
China

Email: peng.shaofu@zte.com.cn

