

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 3, 2022

H. Li
M. Chen
C. Lin
H3C
August 2, 2021

BGP Extensions of SR Policy for Composite Candidate Path
draft-li-idr-sr-policy-composite-path-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 associated with one or more candidate paths. A candidate path is either dynamic, explicit or composite. This document defines extensions to BGP to distribute SR policies carrying composite candidate path information. So that composite candidate paths can be installed when the SR policy is applied.

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 February 3, 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 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
2.	Terminology	2
3.	Constituent SR Policy Attributes in SR Policy	3
3.1.	Constituent SR Policy Sub-TLV	4
4.	Operations	5
5.	Security Considerations	5
6.	IANA Considerations	5
7.	Contributors	6
8.	References	6
8.1.	Normative References	6
8.2.	Informative References	7
	Authors' Addresses	7

[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 [[I-D.ietf-spring-segment-routing-policy](#)]. In order to distribute SR policies to the headend, [[I-D.ietf-idr-segment-routing-te-policy](#)] specifies a mechanism by using BGP.

An SR Policy is associated with one or more candidate paths. A composite candidate path acts as a container for grouping of SR Policies. As described in section 2.2 in [[I-D.ietf-spring-segment-routing-policy](#)], the composite candidate path construct enables combination of SR Policies, each with explicit candidate paths and/or dynamic candidate paths with potentially different optimization objectives and constraints, for a load-balanced steering of packet flows over its constituent SR Policies.

This document defines extensions to Border Gateway Protocol (BGP) to distribute SR policies carrying composite candidate path information. So that composite candidate paths can be installed when the SR policy is applied.

[2.](#) Terminology

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.

3. Constituent SR Policy Attributes 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

...

...

As described in section 2.2 in [[I-D.ietf-spring-segment-routing-policy](#)], the endpoints of the constituent SR Policies and the parent SR Policy MUST be identical, and the colors of each of the constituent SR Policies and the parent SR Policy MUST be different. Therefore a constituent SR Policy is referenced only by color in the composite candidate path since its headend and endpoint are identical to the parent SR policy.

SR policy with composite candidate path information is 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

Explicit NULL Label Policy (ENLP)

Segment List

Weight

Segment

Segment

...

Constituent SR Policy

Weight

...

3.1. Constituent SR Policy Sub-TLV

The Constituent SR Policy sub-TLV encodes a single composite path towards the endpoint. The Constituent SR Policy sub-TLV is an optional sub-TLV of BGP Tunnel Encapsulation Attribute, and MAY appear multiple times in the SR Policy encoding. The ordering of Constituent SR Policy sub-TLVs does not matter. The Constituent SR Policy sub-TLV MAY contain a Weight sub-TLV.

Since a candidate path is either dynamic, explicit or composite, the Constituent SR Policy sub-TLV and the Segment List sub-TLV SHOULD NOT appear in the same candidate path.

The Constituent SR Policy 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																			
Color										sub-TLVs																													

where:

- o Type: to be assigned by IANA.
- o Length: the total length of the value field not including Type and Length fields.
- o RESERVED: 2 octet of reserved bits. SHOULD be set to zero on transmission and MUST be ignored on receipt.
- o Color: 4-octet value identifying the constituent SR policy.
- o sub-TLVs currently defined:
 - * An optional single Weight sub-TLV which is defined in [section 2.4.4.1](#) in [[I-D.ietf-idr-segment-routing-te-policy](#)]. According to [[I-D.ietf-spring-segment-routing-policy](#)], the fraction of flows steered into each constituent SR Policy is equal to the relative weight of each constituent SR Policy.

4. Operations

The document does not bring new operation beyond the description of operations defined in [[I-D.ietf-idr-segment-routing-te-policy](#)]. The existing operations defined in [[I-D.ietf-idr-segment-routing-te-policy](#)] can apply to this document directly.

Typically but not limit to, the SR policies carrying composite candidate path information are configured by a controller.

After configuration, the SR policies carrying path composite candidate path information will be advertised by BGP update messages. The operation of advertisement is the same as defined in [[I-D.ietf-idr-segment-routing-te-policy](#)], as well as the reception.

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](#)].

6. IANA Considerations

This document defines a new Sub-TLV in registries "SR Policy List Sub-TLVs" [[I-D.ietf-idr-segment-routing-te-policy](#)]:

Value	Description	Reference
TBA	Constituent SR Policy Sub-TLV	This document

7. Contributors

In addition to the authors listed on the front page, the following co-authors have also contributed to this document:

Yuanxiang Qiu
H3C
qiuyuanxiang@h3c.com

Liping Yang
H3C
liping_yang@h3c.com

Yang Wang
H3C
wang.a.yang@h3c.com

8. References

8.1. Normative References

- [I-D.ietf-idr-segment-routing-te-policy]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., Rosen, E., Jain, D., and S. Lin, "Advertising Segment Routing Policies in BGP", [draft-ietf-idr-segment-routing-te-policy-13](#) (work in progress), June 2021.
- [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>>.

8.2. Informative References

[I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Talaulikar, K., Voyer, D., Bogdanov, A., and
P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-13](#) (work in progress),
May 2021.

Authors' Addresses

Hao Li
H3C

Email: lihao@h3c.com

Mengxiao Chen
H3C

Email: chen.mengxiao@h3c.com

Changwang Lin
H3C

Email: linchangwang.04414@h3c.com

