

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

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

**Signaling Composite Candidate Path of SR Policy using BGP-LS
draft-li-idr-bgpls-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, and each candidate path is either dynamic, explicit or composite. This document specifies the extensions to BGP Link State (BGP-LS) to carry composite candidate path information in the advertisement of an 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). 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	3
3.	BGP-LS Extensions for Composite Candidate Path	3
3.1.	Constituent SR Policy TLV	3
4.	Operations	4
5.	Security Considerations	4
6.	IANA Considerations	4
7.	Contributors	5
8.	References	5
8.1.	Normative References	5
8.2.	Informative References	6
	Authors' Addresses	6

[1.](#) Introduction

As described in [[RFC7752](#)], BGP Link State (BGP-LS) provides a mechanism by which link-state and TE information can be collected from networks and shared with external components using the BGP routing protocol.

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](#)].

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.

[[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-LS updates. This document extends it to provide some extra information to carry composite candidate path information in the BGP-LS advertisement.

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. BGP-LS Extensions for Composite Candidate Path

[[RFC7752](#)] defines the BGP-LS NLRI that can be a Node NLRI, a Link NLRI or a Prefix NLRI. The corresponding BGP-LS attribute is a Node Attribute, a Link Attribute or a Prefix Attribute.

[[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 section defines a new sub-TLV which is carried in the optional non-transitive BGP Attribute "LINK_STATE Attribute" defined in [[RFC7752](#)].

3.1. Constituent SR Policy TLV

Segment Routing Policy (SR Policy) architecture is specified in [[I-D.ietf-spring-segment-routing-policy](#)]. A SR Policy can comprise of one or more candidate paths, and each candidate path is either dynamic, explicit or composite. A composite candidate path can comprise of one or more constituent SR policies. 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.

The Constituent SR Policy TLV is used to report the constituent SR policy(s) of a composite candidate path. The TLV has 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                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               RESERVED                           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Color                             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Weight                             |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Sub-TLVs (variable)                //
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


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: 32 bits reserved and MUST be set to 0 on transmission and MUST be ignored on receipt.
- o Color: 4 octets that indicates the color of the constituent SR Policy.
- o Weight: 4 octet field that indicates the weight associated with the SID-List for weighted load-balancing. Refer [Section 2.2](#) and 2.11 of [[I-D.ietf-spring-segment-routing-policy](#)].
- o Sub-TLVs: no sub-TLV is currently defined.

4. Operations

The document does not bring new operation beyond the description of operations defined in [[RFC7752](#)] and [[I-D.ietf-idr-te-lsp-distribution](#)]. The existing operations defined in [[RFC7752](#)] and [[I-D.ietf-idr-te-lsp-distribution](#)] can apply to this document directly.

Typically but not limit to, the BGP-LS messages carrying composite candidate path information along with the SR policy are distributed to a controller.

After configuration, the composite candidate path information will be advertised by BGP update messages. The operation of advertisement is the same as defined in [[RFC7752](#)] and [[I-D.ietf-idr-te-lsp-distribution](#)], 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-te-lsp-distribution](#)].

6. IANA Considerations

This document defines a new TLV in the BGP-LS Link Descriptor and Attribute TLVs:

+-----+	-----+	-----+
Value	Description	Reference
+-----+	-----+	-----+
TBA	Constituent SR Policy 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-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", [draft-ietf-idr-te-lsp-distribution-15](#) (work in progress), May 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),
June 2021.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and
S. Ray, "North-Bound Distribution of Link-State and
Traffic Engineering (TE) Information Using BGP", [RFC 7752](#),
DOI 10.17487/RFC7752, March 2016,
<<https://www.rfc-editor.org/info/rfc7752>>.

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

