Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 26 October 2026

C. Lin
New H3C Technologies
W. Cheng
China Mobile
Z. Ali
A. MahendraBabu
Cisco Systems, Inc
R. Chen
ZTE Corporation

24 April 2026

Signaling Composite Candidate Path of SR Policy using BGP-LS
draft-li-idr-bgpls-sr-policy-composite-path-10

Abstract

SR Policy Architecture [RFC9256] defines the concept of a Composite Candidate Path. A regular SR Policy Candidate Path outputs traffic to a set of Segment Lists, while an SR Policy Composite Candidate Path outputs traffic recursively to a set of SR Policies on the same headend. 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 26 October 2026.

Copyright Notice

Table of Contents

1. Introduction

As described in [RFC9552], 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 [RFC9256].

SR Policy Architecture [RFC9256] defines the concept of a Composite Candidate Path. A regular SR Policy Candidate Path outputs traffic to a set of Segment Lists, while an SR Policy Composite Candidate Path outputs traffic recursively to a set of SR Policies on the same headend.

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 [RFC9256], the composite candidate path construct enables combination of SR Policies, for a load-balanced steering of packet flows over its constituent SR Policies.

[I-D.jiang-idr-sr-policy-composite-path] defines extensions for BGP to distribute SR policies carrying composite candidate path information. While as defined in Section 3.6 of [I-D.ietf-pce-multipath], PCEP signals the Composite Candidate Path.

[RFC9857] 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

[RFC9552] 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. [RFC9857] 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 [RFC9552].

3.1. Composite Candidate Path TLV

Segment Routing Policy (SR Policy) architecture is specified in [RFC9256]. 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 Composite Candidate Path TLV is used to report the constituent SR policy(s) of a composite candidate path. It is carried in the optional non-transitive BGP- LS Attribute defined in [RFC9552] and is associated with the SR Policy Candidate Path NLRI type. Only a single instance of this TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.The TLV has following format:

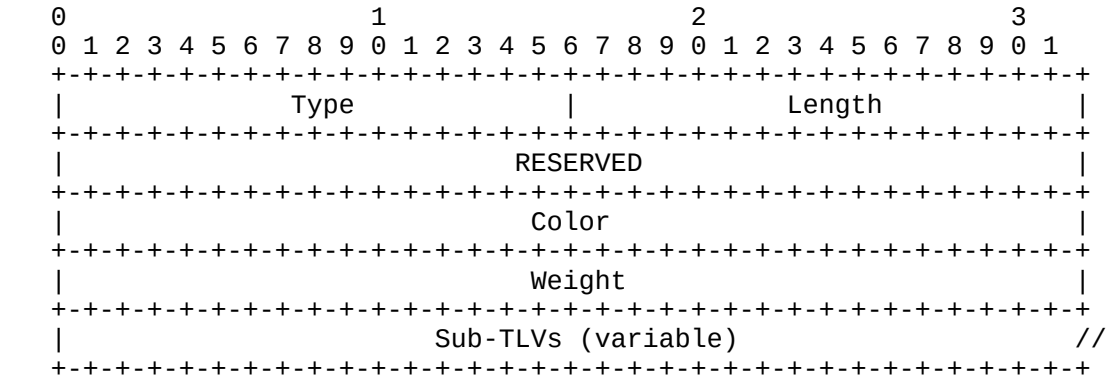

where:

* Type: to be assigned by IANA.

* Length: the total length of the value field not including Type and Length fields.

* Reserved: 32 bits reserved and MUST be set to 0 on transmission and MUST be ignored on receipt.

* Color: 4 octets that indicates the color of the constituent SR Policy.

* 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 [RFC9256].

* Sub-TLVs: variable and contains any other optional attributes associated with the Composite Candidate Path. Currently, the sub-TLV only defines the Per-Flow Forwarding Class TLV.

3.2. Per-Flow Forwarding Class TLV

Per-Flow Candidate Path builds on top of the concept of the Composite Candidate Path. Each Path in a Per-Flow Candidate Path is assigned a 3-bit forward class value, which allows Quality of Service (QoS) classified traffic to be steered depending on the forward class. The Per-Flow Forwarding Class TLV is an optional sub-TLV of the Composite Candidate Path TLV.Only a single instance of this sub-TLV is advertised for a given candidate path. If multiple instances are present, then the first valid (i.e., not determined to be malformed as per section 8.2.2 of [RFC9552]) one is used and the rest are ignored.

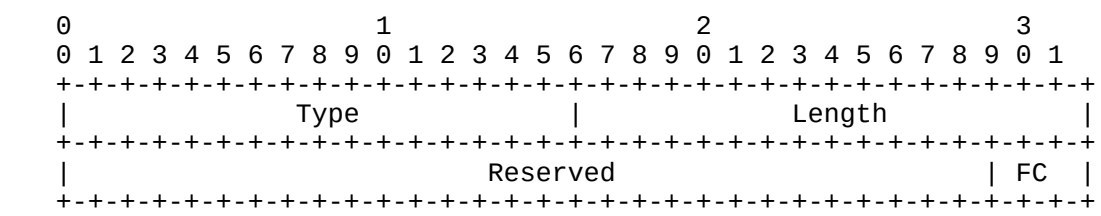

where:

* Type (16 bits): TBD1 for "PER-FLOW-FORWARD-CLASS" TLV.

* Length (16 bits): 4.

* Reserved: This field MUST be set to zero on transmission and MUST be ignored on receipt.

* FC (3 bits): Forward class value that is given by the QoS classifier to traffic entering the given Candidate Path. Different classes of traffic that enter the given Candidate Path can be differentially steered into different Colors.

4. Operations

The document does not bring new operation beyond the description of operations defined in [RFC9552] and [RFC9857]. The existing operations defined in [RFC9552] and [RFC9857] 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 [RFC9552] and [RFC9857], as well as the reception.

5. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the "Security Considerations" section of [RFC4271] for a discussion of BGP security. Security considerations for acquiring and distributing BGP-LS information are discussed in [RFC9552]. Security considerations for acquiring and distributing BGP-LS SR Policy information are discussed in [RFC9857].

Additionally, reporting SR policies that carry composite candidate path information MAY pose a risk to the confidentiality of mission-critical or commercially sensitive network information. It is the responsibility of the network operator to ensure that only trusted nodes (including both routers and controller applications) within the SR domain are configured to receive such information.

6. IANA Considerations

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

Value	Description	Reference
TBA	Composite Candidate Path	This document
TBA	Per-Flow Forwarding Class	This document

Table 1

7. References

7.1. Normative References

[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>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[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>.

[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>.

[RFC9857] Previdi, S., Talaulikar, K., Ed., Dong, J., Gredler, H., and J. Tantsura, "Advertisement of Segment Routing Policies Using BGP - Link State", RFC 9857, DOI 10.17487/RFC9857, October 2025, <https://www.rfc-editor.org/info/rfc9857>.

7.2. Informative References

[I-D.ietf-pce-multipath] Koldychev, M. and S. Sidor, "Path Computation Element Communication Protocol (PCEP) Extensions for Signaling Multipath Information", Work in Progress, Internet-Draft, draft-ietf-pce-multipath-23, 1 April 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-pce-multipath-23>.

[I-D.ietf-spring-sr-policy-group] Cheng, W., Wenying, J., Lin, C., Chen, R., and Y. Zhang, "SR Policy Group", Work in Progress, Internet-Draft, draft-ietf-spring-sr-policy-group-00, 13 January 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-sr-policy-group-00>.

[I-D.jiang-idr-sr-policy-composite-path] Jiang, W., Lin, C., and R. Chen, "BGP Extensions of SR Policy for Composite Candidate Path", Work in Progress, Internet-Draft, draft-jiang-idr-sr-policy-composite-path-07, 23 April 2026, <https://datatracker.ietf.org/doc/html/draft-jiang-idr-sr-policy-composite-path-07>.

[RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <https://www.rfc-editor.org/info/rfc9552>.

Appendix A. Cross-WG Information

This section describes cross-working group information for use during the IETF review process. This section will be removed by the RFC editor prior to publication.

A.1. Spring WG

Section 2.2 of [RFC9256] defines the concept of a Composite Candidate Path. Section 2.13 of [RFC9256] illustrates an information model for hierarchical relationships between the SR Policy constructs. A Parent SR Policy is the composite candidate path that acts as a container for grouping SR Policies which meet different service optimization objectives and constraints and have the same destination endpoint. [I-D. draft-ietf-spring-sr-policy-group] defines illustrates some use cases for parent SR Policy and SR Policy Group to simplify deployment and provide best practice cases for operators.

A.2. PCE WG

Composite Candidate Paths can be distributed via the Path Computation Element Communication Protocol (PCEP), as described in section 3.6 of [I-D.draft-ietf-pce-multipath].

A.3. SRv6Ops

[I-D.jiang-idr-sr-policy-composite-path] defines extensions for BGP to distribute SR policies carrying composite candidate path information. This document extends it to provide some extra information to carry composite candidate path information in the BGP-LS advertisement.

Authors' Addresses

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

Weiqiang Cheng
China Mobile
Email: chengweiqiang@chinamobile.com

Zafar Ali
Cisco Systems, Inc
Email: zali@cisco.com

Aravind Babu MahendraBabu
Cisco Systems, Inc
Email: aramahen@cisco.com

Ran Chen
ZTE Corporation
Nanjing
China
Email: chen.ran@zte.com.cn