

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 9, 2020

Z. Li
L. Li
Huawei
H. Chen
Futurewei
Y. Fan
Casa Systems
X. Liu
Volta Networks
L. Liu
Fujitsu
March 8, 2020

BGP Request for Candidate Paths of SR TE Policies
draft-li-ldr-bgp-request-cp-sr-te-policy-01

Abstract

An SR Policy is a set of candidate paths. The headend of an SR Policy may learn multiple candidate paths for an SR Policy via a number of different mechanisms, e.g., CLI, NetConf, PCEP, or BGP. This document defines extensions to BGP for the headend to request BGP speaker (controller) for advertising the candidate paths.

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 [RFC 2119](#) [[RFC2119](#)].

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 September 9, 2020.

Copyright Notice

Copyright (c) 2020 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. | Overview of BGP Request for SR-TE Paths | 3 |
| 4. | BGP Request UPDATE Message | 4 |
| 4.1. | Extention of SR Policy NLRI | 4 |
| 4.2. | New SR Policy Sub-TLVs | 5 |
| 4.2.1. | SR Path Attributes Sub-TLV | 5 |
| 4.2.2. | Synchronization Sub-TLV | 6 |
| 4.2.3. | Metric Sub-TLV | 8 |
| 4.2.4. | Include Route Sub-TLV | 9 |
| 4.2.5. | Load Balance Sub-TLV | 10 |
| 4.2.6. | Request Parameter Sub-TLV | 10 |
| 5. | IANA | 11 |
| 6. | Contributors | 12 |
| 7. | Acknowledgments | 12 |
| 8. | References | 12 |
| 8.1. | Normative References | 12 |
| 8.2. | Informative References | 13 |
| | Authors' Addresses | 14 |

[1.](#) Introduction

An SR Policy defined in [[I-D.ietf-spring-segment-routing-policy](#)] is a set of candidate paths. The headend of an SR Policy may be informed by various means including: Configuration, PCEP [[RFC8281](#)] or BGP [[I-D.ietf-idr-segment-routing-te-policy](#)]. All these mechanisms are Controller initiated, but in some situations the headend may want to pull a set of candidate paths from Controller rather than receive all information passively. Actually PCEP can use request and reply messages defined in [[RFC5440](#)] to match this requirement, but the

mechanism is not clear when controller advertises candidate paths via BGP.

This document defines a way to request controller (BGP speaker) to advertise candidate paths via BGP update messages. This makes BGP have the mechanism with request and reply similar to PCEP.

2. Terminology

RP: Request Parameters

LSPA: LSP Attributes

SVEC: Synchronization VECTOR

IRO: Include Route Object

ERO: Explicit Route Object

MSD: Base MPLS Imposition Maximum SID Depth, as defined in [[RFC8491](#)]

NAI: Node or Adjacency Identifier

PCC: Path Computation Client

PCE: Path Computation Element

PCEP: Path Computation Element Communication Protocol

SID: Segment Identifier

SR: Segment Routing

SR-TE: Segment Routing Traffic Engineering

3. Overview of BGP Request for SR-TE Paths

[I-D.ietf-idr-segment-routing-te-policy] defines the extensions to BGP for a headend to receive candidate paths in a BGP UPDATE message from a controller (BGP speaker). In some situations a headend just wants to get these candidate paths when some special event occurs (for example, when it receives a customer route (VPN route) with a special color or special BGP attribute). This document defines the mechanism in which the headend requests the controller to advertise the expected SR policy with the candidate paths when this special situation occurs.

At first, the headend decides to get a new candidate path from the controller based on some trigger event. This trigger mechanism is out of scope of this document.

Then, the headend creates a new BGP request UPDATE message (defined below in this document) and sends it to the controller. The message contains the constraints/attributes of SR-TE paths such as affinity, metric, SRLG, and so on. This special request UPDATE message is called request message or request for short. It SHOULD NOT be used for BGP best path selection.

After receiving the request message, the controller will calculate one or a set of paths (i.e., segment lists) according to the request from the headend and advertise the SR Policy with the paths computed to the headend using [[I-D.ietf-idr-segment-routing-te-policy](#)]. How to calculate the paths is out of scope of this document.

4. BGP Request UPDATE Message

A BGP request UPDATE message is based on the update message defined in [[I-D.ietf-idr-segment-routing-te-policy](#)] with some extensions described below.

4.1. Extention of SR Policy NLRI

The SR Policy NLRI defined in [[I-D.ietf-idr-segment-routing-te-policy](#)] has the following format:

```
+-----+
|  NLRI Length   | 1 octet
+-----+
|  Distinguisher  | 4 octets
+-----+
|  Policy Color   | 4 octets
+-----+
|  Endpoint       | 4 or 16 octets
+-----+
```

where:

- o NLRI Length: 1 octet of length expressed in bits as defined in [[RFC4760](#)].
- o Distinguisher: 4-octet value uniquely identifying the policy in the context of <color, endpoint> tuple. The distinguisher has no semantic value and is solely used by the SR Policy originator to make unique (from an NLRI perspective) multiple occurrences of the same SR Policy.

- o Policy Color: 4-octet value identifying (with the endpoint) the policy. The color is used to match the color of the destination prefixes to steer traffic into the SR Policy
[[I-D.ietf-spring-segment-routing-policy](#)]
- o Endpoint: identifies the endpoint of a policy. The Endpoint may represent a single node or a set of nodes (e.g., an anycast address). The Endpoint is an IPv4 (4-octet) address or an IPv6 (16-octet) address according to the AFI of the NLRI.

NLRI Length, Policy Color, Endpoint field remains unchanged, while the Distinguisher field will be set to FF:FF:FF:FF to indicate that the UPDATE message with this NLRI is a request message to the controller.

[4.2.](#) New SR Policy Sub-TLVs

The content of the SR Policy is encoded in the Tunnel Encapsulation Attribute TLV of type 23 defined in [[I-D.ietf-idr-tunnel-encaps](#)] containing a new Tunnel Type TLV of type 15. The SR Policy Encoding structure is as follows:

SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>

Attributes:

Tunnel Encaps Attribute (23)
Tunnel Type (15): SR Policy
 <Sub-TLVs>

Preference, Binding SID, Priority, Policy Name, ENLP, Segment List, Weight and Segment Sub-TLVs are all defined in [[I-D.ietf-idr-segment-routing-te-policy](#)] for a SR Policy to be advertised to a headend.

Additional 6 new Sub-TLVs are defined below for the request mechanism. They are SR Path Attributes, Synchronization, Metric, Include Route, Load Balance, and Request Parameters Sub-TLVs.

[4.2.1.](#) SR Path Attributes Sub-TLV

A SR Path Attributes Sub-TLV contains the attributes of the SR paths requested, which are similar to an LSP Attributes Object defined in [[RFC5440](#)] and [[RFC3209](#)].

It is optional and specifies various attributes or constraints of the paths requested. Its format is shown below.

| 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 | | | | | | | | | | Flags | | | | | | | | | | Reserved | | | | | | | | | |
| + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | | | | | | | | |
| | | | | | | | | | | Exclude-any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | | | | | | | | |
| | | | | | | | | | | Include-any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | | | | | | | | |
| | | | | | | | | | | Include-all | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | | | | | | | | |
| | | | | | | | | | | Optional sub-TLVs | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | + | - | | | | | | | | |

- o Type: TBD3.
- o Length: specifies the length of the value field not including Type and Length fields.
- o Flags (8 bits): Two flags are currently defined:
 - * B (Bound - 1 bit): When set, the metric-value indicates a bound (a maximum) for the path metric that must not be exceeded for the headend to consider the computed path as acceptable. The path metric must be less than or equal to the value specified in the metric-value field. When the B flag is cleared, the metric-value field is not used to reflect a bound constraint.
 - * C (Computed Metric - 1 bit): When set, it indicates that the controller MUST provide the computed path metric value (should a path satisfying the constraints be found) in the advertisement message for the corresponding metric.
 - * Unassigned flags MUST be set to zero on transmission and MUST be ignored on receipt.
- o T (Type - 8 bits): Specifies the metric type. Four metric types are currently defined:
 - * T=1: IGP metric
 - * T=2: TE metric
 - * T=3: Hop Counts
 - * T=11: Maximum SID Depth of the requested path

- o Metric-Value (32 bits): It is a metric value encoded in 32 bits IEEE floating point format (see [IEEE.754.1985]).

4.2.4. Include Route Sub-TLV

The Include Route Sub-TLV is optional and can be used to specify that the computed candidate path MUST traverse a set of specified network elements. The Include Route Sub-TLV carries the same content as IRO Object defined in[RFC5440], [RFC3209] and SR-ERO defined in [I-D.ietf-pce-segment-routing]

The Include Route Sub-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      |      Length      |   NT   |      Flags   |F|S|C|M|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               SID (optional)                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               NAI (variable, optional)                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Where:

- o Type: TBD4.
- o Length: It specifies the length of the value field not including Type and Length fields.
- o NAI Type (NT): It indicates the type and format of the NAI contained, if any is present. If the F bit is set to zero, then the NT field has no meaning and MUST be ignored by the receiver. This document describes the following NT values:
 - * NT=0: The NAI is absent.
 - * NT=1: The NAI is an IPv4 node ID.
 - * NT=2: The NAI is an IPv6 node ID.
 - * NT=3: The NAI is an IPv4 adjacency.
 - * NT=4: The NAI is an IPv6 adjacency with global IPv6 addresses.
 - * NT=5: The NAI is an unnumbered adjacency with IPv4 node IDs.

- * NT=6: The NAI is an IPv6 adjacency with link-local IPv6 addresses.

- o SID and NAI are the same as SR-ERO defined in [\[I-D.ietf-pce-segment-routing\]](#)

4.2.5. Load Balance Sub-TLV

A Load Balance Sub-TLV specifies how many SR paths (i.e., segment lists) should be computed for a path request. It 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   | Length | Flag | Max-Slist |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               Optional sub-TLVs                               ~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

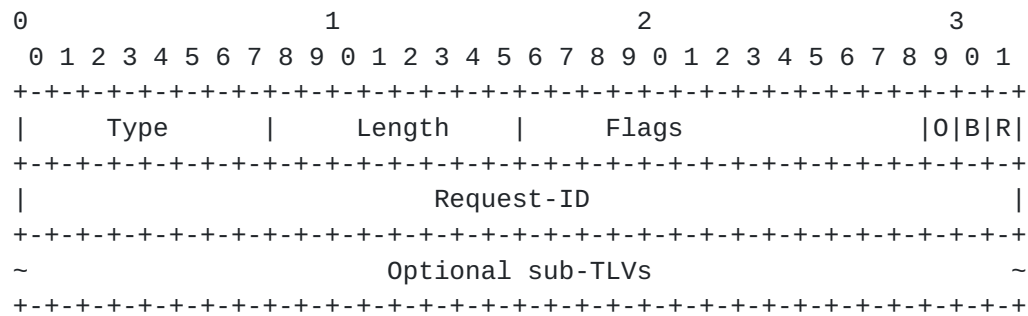
```

Where:

- o Type: TBD5.
- o Length: It specifies the length of the value field not including Type and Length fields.
- o Flags (8 bits): No flag is currently defined. The Flags field MUST be set to zero on transmission and MUST be ignored on receipt.
- o Max-Slist (8 bits): It indicates the maximum number of SR paths (i.e., segment lists) to be computed for the request. The load is distributed among these SR paths.
- o Optional sub-TLVs: No Optional sub-TLV is currently defined.

4.2.6. Request Parameter Sub-TLV

A Request Parameter (RP) Sub-TLV specifies the request identifier and other parameters for a path request. It has the format below.



Where:

- o Type: TBD6.
- o Length: It specifies the length of the value field not including Type and Length fields.
- o Flags (16 bits): Three flag bits are currently defined as follows:
 - * R (Reoptimization - 1 bit): when set, it indicates that the SR path request message is for the reoptimization of an existing SR path, which is represented by a segment list Sub-TLV in the message.
 - * B (Bi-directional - 1 bit): when set, it indicates that the SR path request relates to bi-directional paths that has the same traffic engineering requirements including fate sharing, TE links, and other requirements (such as latency and jitter) in each direction.
 - * O (strict/loose - 1 bit): when set, it indicates that a loose path is acceptable. Otherwise (i.e., when cleared), it indicates that a path exclusively made of strict hops is required.

5. IANA

Under Existing Registry Name: "BGP Tunnel Encapsulation Attribute Sub-TLVs", IANA is requested to assign new Sub-TLV values for SR Path Request as follows:

| Type | Value | Sub-TLV Name | Reference |
|------|-------|----------------------------|---------------|
| TBD1 | | SR Path Attributes Sub-TLV | This document |
| TBD2 | | Synchronization Sub-TLV | This document |
| TBD3 | | Metric Sub-TLV | This document |
| TBD4 | | Include Route Sub-TLV | This document |
| TBD5 | | Load Balance Sub-TLV | This document |
| TBD6 | | Request Parameters Sub-TLV | This document |

6. Contributors

TBD

7. Acknowledgments

TBD

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-08](#) (work in progress), November 2019.
- [I-D.ietf-spring-segment-routing-policy]
 Filsfils, C., Sivabalan, S., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-06](#) (work in progress), December 2019.
- [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>>.

- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), DOI 10.17487/RFC5440, March 2009, <<https://www.rfc-editor.org/info/rfc5440>>.
- [RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", [RFC 8281](#), DOI 10.17487/RFC8281, December 2017, <<https://www.rfc-editor.org/info/rfc8281>>.

8.2. Informative References

- [I-D.ietf-idr-tunnel-encaps]
Patel, K., Velde, G., and S. Ramachandra, "The BGP Tunnel Encapsulation Attribute", [draft-ietf-idr-tunnel-encaps-15](#) (work in progress), December 2019.
- [I-D.ietf-pce-segment-routing]
Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "PCEP Extensions for Segment Routing", [draft-ietf-pce-segment-routing-16](#) (work in progress), March 2019.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", [RFC 4090](#), DOI 10.17487/RFC4090, May 2005, <<https://www.rfc-editor.org/info/rfc4090>>.
- [RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", [RFC 4760](#), DOI 10.17487/RFC4760, January 2007, <<https://www.rfc-editor.org/info/rfc4760>>.
- [RFC5420] Farrel, A., Ed., Papadimitriou, D., Vasseur, JP., and A. Ayyangarps, "Encoding of Attributes for MPLS LSP Establishment Using Resource Reservation Protocol Traffic Engineering (RSVP-TE)", [RFC 5420](#), DOI 10.17487/RFC5420, February 2009, <<https://www.rfc-editor.org/info/rfc5420>>.

Authors' Addresses

Zhenbin Li
Huawei
156 Beiqing Road
Beijing, 100095
P.R. China

Email: lizhenbin@huawei.com

Lei Li
Huawei
156 Beiqing Road
Beijing, 100095
P.R. China

Email: lily.lilei@huawei.com

Huaimo Chen
Futurewei
Boston, MA
USA

Email: Huaimo.chen@futurewei.com

Yanhe Fan
Casa Systems
USA

Email: yfan@casa-systems.com

Xufeng Liu
Volta Networks
McLean, VA
USA

Email: xufeng.liu.ietf@gmail.com

Lei Liu
Fujitsu
USA

Email: liulei.kddi@gmail.com

