

Inter-Domain Routing Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 2, 2021

Y. Wang
H. Li
Y. Qiu
L. Yang
M. Chen
H3C Technologies
February 2, 2021

Segment Routing PCE Delegation in BGP
draft-wang-idr-sr-policy-pce-delegation-00

Abstract

Segment Routing is a source routing paradigm that explicitly indicates the forwarding path for packets at the ingress node. An SR policy is a set of candidate SR paths consisting of one or more segment lists with necessary path attributes. The headend of an SR Policy may learn multiple candidate paths for an SR Policy. Candidate paths may be learned via a number of different mechanisms, e.g., CLI, NetConf, PCEP, or BGP.

The Path Computation Element (PCE) provides path computation functions in support of traffic engineering in Multi Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks.

Currently, when a controller uses BGP to deploy an SR Policy, there is no way to encode PCE delegation related options. The only way to import a PCE delegation is through local configuration management, e.g., Netconf, CLI or gRPC.

This document defines extensions to BGP to distribute PCE delegation information within 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 August 2, 2021.

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	3
2.	Terminology	3
2.1.	Requirements Language	4
3.	SR Policy for PCE Delegation.	4
3.1.	PCE Delegation Sub-TLV	6
4.	Operations	7
5.	IANA Considerations	8
6.	Security Considerations	8
7.	Contributors	9
8.	Acknowledgements	9
9.	References	9
9.1.	Normative References	9
	Authors' Addresses	10

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.

[RFC5440] describes the Path Computation Element (PCE) Communication Protocol (PCEP). PCEP enables the communication between a Path Computation Client (PCC) and a PCE, or between PCE and PCE, for the purpose of computation of Multi protocol Label Switching (MPLS) as well as Generalized MPLS (GMPLS) Traffic Engineering Label Switched Path (TE LSP) characteristics.

When a controller uses BGP to deploy an SR Policy, the only way to import a PCE delegation is through configuration management e.g., Netconf, CLI or gRPC. It's inconvenient for a controller to use separate channels to deploy SR policies.

This document defines extensions to BGP to distribute PCE delegation information within SR policies.

2. Terminology

PCE delegation : An operation to grant a PCE temporary rights to modify a subset of LSP parameters on one or more PCC's LSPs. LSPs are delegated from a PCC to a PCE, and are referred to as delegated LSPs. The PCC who owns the PCE state for the LSP has the right to delegate it. An LSP is owned by a single PCC at any given point in time. For intra-domain LSPs, this PCC SHOULD be the PCC of the LSP head end.

2.1. Requirements Language

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. SR Policy for PCE Delegation

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
        ...
    ...
```

As introduced in [Section 1](#), SR Policy could use PCE to calculate path. SR policy with PCE delegation information is expressed as below:

```
SR Policy SAFI NLRI: <Distinguisher, Policy-Color, Endpoint>
Attributes:
  Tunnel Encaps Attribute (23)
    Tunnel Type: SR Policy
      Binding SID
      Preference
      Priority
      Policy Name
      PCE Delegation
      Explicit NULL Label Policy (ENLP)
      Segment List
        Weight
        Segment
        Segment
        ...
      ...
```

3.1. PCE Delegation Sub-TLV

A PCE Delegation sub-TLV is an Optional sub-TLV. When it appears, it must appear only once at most within a SR Policy Sub-TLV. If multiple PCE Delegation sub-TLVs appear within an SR Policy Sub-TLV, the NLRI MUST be treated as a malformed NLRI.

As per [[I-D.ietf-idr-segment-routing-te-policy](#)], when the error determined allows for the router to skip the malformed NLRI(s) and continue processing of the rest of the update message, then it MUST handle such malformed NLRIs as 'Treat-as-withdraw'. This document does not define new error handling rules for PCE Delegation sub-TLV, and the error handling rules defined in [[I-D.ietf-idr-segment-routing-te-policy](#)] apply to this document.

PCE Delegation is a new sub-TLV of the BGP Tunnel Encapsulation Attribute [[I-D.ietf-idr-tunnel-encaps](#)].

A PCE Delegation sub-TLV is associated with an SR policy. The PCE Delegation sub-TLV has the following format:





Figure 1. PCE Delegation sub-TLV

Where:

Type: to be assigned by IANA.

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

Flags: 1 octet of flags.



where:

* D-Flag: The flag encode PCE delegation. When this flag is set, it indicates to enable PCE delegation function, otherwise it indicates to disable PCE delegation.

* R-Flag: The flag encode passive delegation report only. An implementation SHOULD report its the status of SR policies to PCE server without using PCE Server to calculate path.

To avoid confusion, D-Flag and R-Flag SHOULD NOT be set simultaneously.

Unused bits in the Flag octet SHOULD be set to zero upon transmission and MUST be ignored upon receipt.

Reserved: 8 bits reserved and MUST be set to 0 on transmission and MUST be ignored on receipt.

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 PCE Delegation information are configured by a controller.

After configuration, the SR policies carrying PCE Delegation 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.

The consumer of the SR policies is not the BGP process. The operation of sending information to consumers is out of scope of this document.

5. 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	PCE Delegation sub-TLV	This document

6. Security Considerations

TBA

7. Contributors

Yang.Wang

H3C Technology

China

Email: wang.a.yang@h3c.com

8. Acknowledgements

Authors would like to thank Changwang.Lin, Jinrong.Ye for their professional comments and help.

9. References

9.1. Normative References

- [I-D.ietf-idr-tunnel-encaps]
Patel, K., Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", [draft-ietf-idr-tunnel-encaps-22](#) (work in progress), January 2021.
- [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-11](#) (work in progress), May 2020.
- [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-09](#) (work in progress), July 2020.
- [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>>.
- [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>>.

Authors' Addresses

Yang Wang
H3C Technologies
China

Email: wang.a.yang@h3c.com

Hao Li
H3C Technologies
China

Email: lihao@h3c.com

Yuanxiang Qiu
H3C Technologies
China

Email: qiuyuanxiang@h3c.com

Liping Yang
H3C Technologies
China

Email: liping_yang@h3c.com

Mengxiao Chen
H3C Technologies
China

Email: chen.mengxiao@h3c.com

