PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: 8 January 2026

Z. Ali
Cisco Systems, Inc.
C. Lin
New H3C Technologies
Y. Liu
China Mobile
R. Chen
ZTE Corporation
C. Li
Huawei Technologies
7 July 2025

Path Computation Element Communication Protocol (PCEP) extensions for SRv6 Policy SID List Optimization
draft-all-pce-srv6-policy-sid-list-optimization-02

Abstract

In some use cases, an SRv6 policy's SID list ends with the policy endpoint's node SID, and the traffic steered (over policy) already ensures that it is taken to the policy endpoint. In such cases, the SID list can be optimized by excluding the endpoint Node SID when installing the policy. This draft specifies a PCEP extension to indicate whether the endpoint's node SID needs to be included or excluded when installing the SRv6 Policy.

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.

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/.

Copyright Notice

Table of Contents

1. Introduction

Segment Routing (SR) [RFC8402] allows a node to steer a packet flow along any path. A Segment Routing Policy (SR Policy) [RFC8402] is an ordered list of segments that represent a source-routed policy. The headend node is said to steer a flow into an SR Policy. The packets steered into an SR Policy have an ordered list of segments associated with that SR Policy written into them. Segment Routing Policy Architecture [RFC9256] updates [RFC8402] as it details the concepts of SR Policy and steering into an SR Policy. [RFC8986] describes the representation and processing of this ordered list of segments for Segment Routing over IPv6 (SRv6). [RFC9603] specifies PCEP extensions to support SR for the IPv6 data plane.

A PCE computes the SRv6 TE Policy SID list from the headend to the endpoint. The computed SID list may end with the policy endpoint’s Node SID or the penultimate hop adjacency SID. If the computed SID list ends with the policy endpoint’s node SID and the overlay SID in the steered traffic (over policy) already ensures that the traffic is taken to the policy endpoint with the same intent, the SRv6 policy endpoint device needs to process back-to-back local node SIDs. Examples of overlay SID containing the local node SID are a service SID, a binding SID for transit policies, an EPE SID, etc. From a compression efficiency viewpoint, carrying back-to-back end-point node SID is inefficient. The SID list in the packet can be optimized by excluding the end-point node SID when installing the policy. End-point node SID exclusion improves the compression efficiency and makes packet processing more efficient for the policy endpoint.

Excluding the policy endpoint’s node SID is possible in most use cases, but not all. For example, if the SRv6 policy is used to carry MPLS traffic, as described in [I-D.draft-agrawal-spring-srv6-mpls-interworking], it is not possible to exclude the policy endpoint’s node SID. Specifically, the endpoint’s node SID inclusion or exclusion is a policy attribute. This draft specifies a PCEP extension to include or exclude the node SID when installing the SRv6 Policy.

The Endpoint Node Suppression (ENS) procedure specified in this draft are equally applicable to PCE initiated LSPs as well as PCC inited LSPs.

2. 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. Terminology

This document uses the following terms defined in [RFC5440]: PCC, PCE, PCEP, PCEP Peer.

SR: Segment Routing.

SID: Segment Identifier.

SRv6: Segment Routing over IPv6 data plane.

4. Overview of PCEP Extensions

4.1. New TLV in the SRPA Object

The draft specifies a new SRPOLICY-POL-ATTRIBUTE TLV for the SR Policy Association object defined in [I-D.draft-ietf-pce-segment-routing-policy-cp]. The SRPOLICY-POL-ATTRIBUTE TLV is optional.

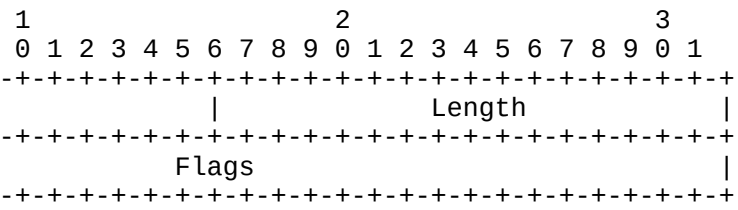

Figure 1: SRPOLICY-POL-ATTRIBUTE TLV

* Type: TBD1 for "SRPOLICY-POL-ATTRIBUTE" TLV.

* Length: 4.

* Flags: The unassigned bits in the Flag octet MUST be set to zero upon transmission and MUST be ignored upon receipt.

4.2. A New flag in the SRPOLICY-POL-ATTRIBUTE TLV

This document specifies ENS-flag (Endpoint Node Suppression flag) bit in the Flags field of "SRPOLICY-POL-ATTRIBUTE" TLV specified in this document. The flag is applicable only to SR policies with SRv6 data plane. The flag MUST be ignored for SR policies with SR-MPLS data plane.

ENS (Endpoint Node Suppression) - 1 bit (Bit Position TBD2):

* If set to 1, the endpoint node SID MUST be excluded when installing the SRv6 Policy SID list(s) used to carry the data traffic.

* If set to 0, the endpoint node SID MUST be included when installing the SRv6 Policy SID list(s) used to carry the data traffic.

4.3. New flag in SRv6-PCE-CAPABILITY sub-TLV

ENS (Endpoint Node Suppression capability flag) is proposed in the SRv6-PCE-CAPABILITY sub-TLV defined in [RFC9603]. The bit position for the flag in the SRv6 Capability Flag Field registry is to be defined by IANA.

ENS (Endpoint Node Suppression flag) - 1 bit (Bit Position TBD3):

* If set to 1, it indicates support for the ENS-flag in the SRPOLICY-POL-ATTRIBUTE TLV.

5. Operation

A PCE always computes the SRv6 TE Policy SID list from the headend to the endpoint (node SID).

A PCEP speaker indicates its ability to support ENS-flag in the Flags field of the SRPOLICY-POL-ATTRIBUTE TLV during the PCEP initialization phase by setting the ENS-flag in the SRv6-PCE-CAPABILITY sub-TLV in the Open message.

A PCEP peer indicates the inclusion or exclusion of the endpoint’s Node SID in ENS-flag in the Flags field of the SRPOLICY-POL-ATTRIBUTE TLV.

A PCEP peer MUST NOT set the ENS-flag flag if capability was not advertised by both peers.

If the computed SID list ends with the policy endpoint's Node SID and the traffic steered over policy already ensures that the traffic is taken to the policy endpoint and the PCEP peers are capable of supporting the ENS-flag, the PCE MUST set ENS-flag to 1.

If the computed SID list ends with the policy endpoint's Node SID and the traffic steered over policy does not takes the traffic to the policy endpoint and the PCEP peers are capable of supporting the ENS-flag, the PCE MUST set ENS-flag to 0.

If the computed SID list ends with the penultimate hop adjacency SID, and the PCEP peers are capable of supporting the ENS-flag, the PCE MUST set ENS-flag to 0.

If the PCEP peers are capable of supporting the ENS-flag and the ENS-flag in the Flags field of the SRPOLICY-POL-ATTRIBUTE TLV is set, the PCC MUST exclude the endpoint node SID when installing the SRv6 Policy sid list(s) used to carry data traffic.

If the PCEP peers are capable of supporting the ENS-flag and the ENS-flag in the Flags field of the SRPOLICY-POL-ATTRIBUTE TLV is not set, the PCC MUST include the endpoint node SID when installing the SRv6 Policy sid list(s) used to carry data traffic.

ENS-flag value in the Flags field of the SRPOLICY-POL-ATTRIBUTE TLV MUST NOT change for a given SRv6 Policy Candidate Path during its lifetime.

Local policy at PCC MAY override the ENS-flag.

PCE ignores the ENS-flag received from the PCC when computing the path and computes the SRv6 Policy SID list from the headend to the endpoint. PCE MAY use the ENS-flag value for debugging purposes.

5.1. MSD Consideration

In some cases, the SID list computed by the PCE exceeds the Maximum Stack Depth (MSD) that the headend node is capable of supporting. In such cases, the PCE has to install transit policies to reduce the sid-list to fit within the MSD capability of the headend node. As the SRv6 policy endpoint node suppression reduces the sid-list size, the section describes the MSD consideration related to this draft.

Suppose the size of the SRv6 TE Policy SID list computed by PCE is L. If the PCEP peers are capable of supporting the ENS-flag, and the PCE sets the ENS-flag to 0, the PCE uses the full sid-list length (L) in the headend MSD consideration procedure. If the PCEP peers are capable of supporting the ENS-flag and the PCE sets the ENS-flag to

1, the PCE uses the sid-list length (L-1) in the headend MSD consideration procedure. This is because the endpoint node SID is suppressed. The MSD consideration procedure is outside the scope of this document.

6. Backward compatibility

If at least one PCEP peer is not capable of supporting the ENS-flag, the endpoint Node SID inclusion/exclusion SHOULD be set based on local policy at the PCC.

7. Security Considerations

[RFC8754] defines the notion of an SR domain and use of SRH within the SR domain. Procedures for securing an SR domain are defined the section 5.1 and section 7 of [RFC8754]. This document does not impose any additional security challenges to be considered beyond security threats described in [RFC8754], [RFC8679] and [RFC8986].

8. IANA Considerations

TBA

9. Contributors

The following people have contributed to this document:

Rajesh M Venkateswaran
Cisco Systems, Inc.
Email: rmelarco@cisco.com

Yuanxiang Qiu
New H3C Technologies
Email: qiuyuanxiang@h3c.com

Samuel Sidor
Cisco Systems, Inc.
Email: ssidor@cisco.com

10. References

10.1. 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>.

[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>.

[RFC8679] Shen, Y., Jeganathan, M., Decraene, B., Gredler, H., Michel, C., and H. Chen, "MPLS Egress Protection Framework", RFC 8679, DOI 10.17487/RFC8679, December 2019, <https://www.rfc-editor.org/info/rfc8679>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

[RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <https://www.rfc-editor.org/info/rfc8986>.

[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>.

[RFC9603] Li, C., Ed., Kaladharan, P., Sivabalan, S., Koldychev, M., and Y. Zhu, "Path Computation Element Communication Protocol (PCEP) Extensions for IPv6 Segment Routing", RFC 9603, DOI 10.17487/RFC9603, July 2024, <https://www.rfc-editor.org/info/rfc9603>.

10.2. Informative References

Appendix A. Acknowledgements

The authors would like to thank Ketan Talaulikar and Andrew Stone for the review comments.

Authors' Addresses

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

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

Yisong Liu
China Mobile
Email: liuyisong@chinamobile.com

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

Cheng Li
Huawei Technologies
Email: c.l@huawei.com