PCE
Internet-Draft
Intended status: Standards Track
Expires: 3 September 2026

Z. Ali
Cisco Systems, Inc.
Y. Liu
ZTE Corporation
A. Stone
D. Achaval
Nokia
S. Sidor
Cisco Systems, Inc.
2 March 2026

MSD Consideration in Path Computation Element Communication Protocol (PCEP)
draft-ali-pce-sr-policy-msd-consideration-01

Abstract

Segment Routing (SR) allows a node to steer a packet flow along any path. SR Policy is an ordered list of segments (i.e., instructions) that represent a source-routed policy. The packets steered into an SR Policy carry an ordered list of segments associated with that SR Policy. An SR Policy can be instantiated SR-MPLS and SRv6 data planes.

Maximum SID Depth (MSD) is first introduced for SR-MPLS to indicate the number of SIDs supported by a node or a link on a node. This concept is further extended for SRv6 with more types of MSD. MSD may become one of the limitations that need to be considered when computing an SR-TE path for PCE.

This draft specifies some MSD considerations PCE needs to take into account when computing an SR-TE path.

This Internet-Draft will expire on 3 September 2026.

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. An SR Policy can be instantiated SR-MPLS and SRv6 data planes.

The concept of Maximum SID Depth (MSD) [RFC8491] is first introduced for SR-MPLS to express the number of SIDs supported by a node or a link on a node, and the Base MPLS Imposition MSD is defined to indicate the number of MPLS labels that can be imposed by a router. And the concept is further extended for SRv6 with more types of MSD defined in [RFC9352].

MSD may become one of the limitations that need to be considered when computing an SR-TE path for PCE.

For SR-MPLS, [RFC8664] defines the SR-PCE-CAPABILITY sub-TLV. PCEP speakers use this sub-TLV to exchange information about their SR capability, including MSD, which indicates that a PCC is capable of imposing on a packet. [RFC8664] also specifies MSD considerations PCE needs to take into account when computing the number of SIDs in an SR-TE path. Specifically, it mandates that once an SR-capable PCEP session is established with a non-zero MSD value, the corresponding PCE MUST NOT send SR-TE paths with a number of SIDs exceeding that MSD value.

Similarly, for SRv6, [RFC9603] specifies that a PCE MUST NOT send SRv6 paths that exceed the SRv6 MSD capabilities of the PCC.

However, the limitation of MSD could be loosen to allow one more SID in the SID list that is sent by the PCE in the following scenarios:

* When the first SID in an SR Policy SID list is an adjacency SID attached by the headend, for SR-MPLS, the top adjacency SID is not imposed on the packet, for SRv6, the implementation can also choose not to include the top adjacency SID in the SRH.

* For SRv6, when a reduced SRH [RFC8754] is used, the first segment of the related SR Policy is not imposed in the reduced SRH.

Moreover, not all of the SRv6 MSDs defined in [RFC9352] are about the limitation/capability of the head-end node (i.e, PCC), thus some of these SRv6 MSDs are not always necessary restrictions to be followed when sending an SRv6 path to the PCC.

This document specifies a procedure for optimizing the number of SIDs in an SR-TE path that PCE can compute when the first SID in the SR Policy SID list is not imposed on the packet in the above scenarios.

This document also analyzes the impact of different SRv6 MSDs when PCE sends SR-TE paths to the PCE.

2. Terminology

The following terminology is used in this document:

MSD: Maximum SID Depth

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

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. PCEP Extensions for SR-MPLS

3.1. New flag in SR-PCE-CAPABILITY sub-TLV

This section proposes a new A-flag (Adjacency SID exclusion for MSD consideration flag) in the SR-PCE-CAPABILITY sub-TLV defined in [RFC8664]. The bit position for the flag in the SR PCE Capability Flag Field registry is to be defined by IANA.

A (Adjacency SID exclusion for MSD consideration flag) - 1 bit (Bit Position TBD1):

* If set to 1, it indicates support for the A-flag by the PCEP peer.

3.2. Operation

[RFC8664] mandates that once an SR-capable PCEP session is established with a non-zero MSD value, the corresponding PCE MUST NOT send SR-TE paths with a number of SIDs exceeding that MSD value.

The following procedure MUST only be applied if both the PCE and PCC have advertised support for the capability by setting the A-flag in their respective SR-PCE-CAPABILITY sub-TLVs [RFC8664]. Under these conditions, if the first SID in an SR-MPLS TE path is an adjacency SID, the PCE MUST NOT send SR-TE paths with a number of SIDs exceeding that (MSD+1) value.

4. Considerations for SRv6 MSDs

[RFC9603] defines the SRv6-PCE-CAPABILITY sub-TLV under the PATH-SETUP-TYPE-CAPABILITY TLV in the OPEN object. PCEP speakers use this sub-TLV to exchange information about their SRv6 capability. And the SRv6 MSD information advertised via SRv6-PCE-Capability sub-TLV conveys the SRv6 capabilities of the PCEP speaker.

As in [RFC9603] section 4.1.1, optional (MSD-Type,MSD-Value) pairs are carried in the SRv6-PCE-CAPABILITY sub-TLV, the SRv6 MSD types are as per [RFC9352], i.e, Maximum Segments Left MSD, Maximum End Pop MSD, Maximum H.Encaps MSD, Maximum End D MSD:

* For Maximum H.Encaps MSD, which indicates the maximum number of SIDs that can be added to the segment list of an SRH as part of the "H.Encaps" behavior, if Maximum H.Encaps MSD is n, actually the PCE can send n+1 SIDs when the first SID is not in the SRH (i.e., when the reduced SRH is used or when the PCC does not impose the first adjacency SID).

* For Maximum Segments Left MSD, when reduced SRH is used, it is not affected since Maximum Segments Left MSD indicates the maximum value of the "Segments Left" field [RFC8754] in the SRH, and it is not related with whether the first SID is in the SRH.

* For Maximum End Pop MSD Type, it signals the maximum number of SIDs in the SRH to which the router can apply "Penultimate Segment Pop (PSP)" as the the penultimate SR Segment Endpoint Node or "Ultimate Segment Pop (USP) " as the ultimate SR Segment Endpoint Node, as defined in "Flavors" (Section 4.16 of [RFC8986]). So usually this limitation does not apply for the head-end node(acting as a PCC), unless the head-end nodes is also the penultimate or the ultimate node in the same SID-list.

* For Maximum End D MSD, it specifies the maximum number of SIDs present in an SRH when performing decapsulation(e.g, End.DX6, End.DT4, End.DT46, End with USD, and End.X with USD [RFC8986]). Similar with Maximum End Pop MSD, the head-end node of an SRv6 path normally would not perform decapsulation.

To conclude, when appears in the SRv6-PCE-CAPABILITY sub-TLV, the Maximum End Pop MSD and Maximum End D MSD only indicates the limitation when the PCC node acts as the penultimate or ultimate SR Segment Endpoint Node. So the Maximum End Pop MSD or Maximum End D MSD is not considered by PCE when sending the SRv6 path to PCC/ head-end node. The limitation of Maximum H.Encaps MSD could be loosen when the first SID is not in the SRH.

5. PCEP Extensions for SRv6

5.1. A-Flag in SRv6-PCE-CAPABILITY sub-TLV

This section proposes a new A-flag (Adjacency SID exclusion for MSD consideration flag) 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.

A (Adjacency SID exclusion for MSD consideration flag) - 1 bit (Bit Position TBD1):

* If set to 1, it indicates support for the A-flag by the PCEP peer.

5.2. R-Flag in SRv6-PCE-CAPABILITY sub-TLV

This section proposes an R-Flag (Reduced SRH for MSD consideration flag) 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.

R-flag (Reduced SRH for MSD consideration flag) - 1 bit (Bit Position TBD2):

* If set to 1, it indicates support for the R-flag by the PCEP peer.

5.3. Operation

[RFC9603] specifies that a PCE MUST NOT send SRv6 paths that exceed the SRv6 MSD capabilities of the PCC.

For the A-Flag, when both the PCE and PCC have advertised support for the capability by setting the A-flag in their respective SRv6-PCE-CAPABILITY sub-TLVs [RFC9603], under these conditions, if the first SID in an SRv6-TE path is an adjacency SID attached with the headend node, the PCE MUST NOT send SR-TE paths with a number of SIDs exceeding that (Maximum H.Encaps MSD+1) value.

For the R-Flag, when both the PCE and PCC have advertised support for the capability by setting the R-flag in their respective SRv6-PCE-CAPABILITY sub-TLVs [RFC9603], the PCE MUST NOT send SR-TE paths with a number of SIDs exceeding the (Maximum H.Encaps MSD+1) value.

(To be discussed) If both the A-Flag and R-Flag are set in the respective SRv6-PCE-CAPABILITY sub-TLVs [RFC9603] of PCE and PCC, the PCE MUST NOT send SR-TE paths with a number of SIDs exceeding the (Maximum H.Encaps MSD+1) value, since A-Flag and R-Flag are both about omitting the first SID in the SID list.

6. Backward compatibility

The proposed procedure is backward compatible with [RFC8664] and [RFC9603] as it requires both PCE and PCC to support the optimization capabilities during the PCEP initialization phase by setting the corresponding new flag in the SR-PCE-CAPABILITY or SRv6-PCE-CAPABILITY sub-TLV in the Open message. Specifically, if at least one PCEP peer is not capable of supporting the new flags, the PCE MUST NOT send SR-TE paths with a number of SIDs exceeding that MSD capability.

7. Security Considerations

Security considerations in [RFC8664] and [RFC9603] apply to this document.

8. IANA Considerations

This document requests IANA to assign an R-Flag in the "SRv6 Capability Flag Field" registry.

Bit	Description	Reference
TBD1	Adjacency SID exclusion for MSD consideration (A-Flag)	[this document]

This document requests IANA to assign the following flags in the "SRv6 Capability Flag Field" registry.

Bit	Description	Reference
TBD1	Adjacency SID exclusion for MSD consideration (A-Flag)	[this document]
TBD2	Reduced SRH for MSD consideration(R-Flag)	[this document]

9. References

9.1. 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>.

[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>.

[RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <https://www.rfc-editor.org/info/rfc8664>.

[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>.

9.2. Informative References

[RFC8491] Tantsura, J., Chunduri, U., Aldrin, S., and L. Ginsberg, "Signaling Maximum SID Depth (MSD) Using IS-IS", RFC 8491, DOI 10.17487/RFC8491, November 2018, <https://www.rfc-editor.org/info/rfc8491>.

[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>.

[RFC9352] Psenak, P., Ed., Filsfils, C., Bashandy, A., Decraene, B., and Z. Hu, "IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane", RFC 9352, DOI 10.17487/RFC9352, February 2023, <https://www.rfc-editor.org/info/rfc9352>.

Authors' Addresses

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

Yao Liu
ZTE Corporation
Email: zliu.yao71@zte.com.cn

Andrew Stone
Nokia
Email: andrew.stone@nokia.com

Diego Achaval
Nokia
Email: diego.achaval@nokia.com

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