

PCE
Internet-Draft
Updates: RFC9603 (if approved)
Intended status: Standards Track
Expires: 15 April 2026

Y. Liu
ZTE Corporation
12 October 2025

Considerations for SRv6 MSD Limitation in PCEP
draft-liu-pce-srv6-msd-consideration-00

Abstract

This document analyzes the impact of different SRv6 MSDs when PCE sends SR-TE paths to the PCE. And it updates the MSD restrictions in RFC9603 based on the analysis. This document also specifies a procedure for optimizing the number of SIDs in an SRv6-TE path that PCE can compute when the reduced SRH mode is used.

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 15 April 2026.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
2.1. Requirements Language	3
3. SRv6 MSDs	3
4. Updates to RFC9603	4
5. New Flag in SRv6-PCE-CAPABILITY sub-TLV	5
6. Operational Considerations	5
7. Security Considerations	5
8. IANA Considerations	5
9. References	5
9.1. Normative References	5
9.2. Informative References	6
Author's Address	6

1. Introduction

As defined in [RFC8402], Segment Routing (SR) architecture allows the source node to steer a packet through a path indicated by an ordered list of instructions, called "segments". A segment can represent any instruction, topological or service based, and it can have a semantic local to an SR node or global within an SR domain.

[RFC8664] specifies PCEP extensions for supporting an SR-TE LSP for the MPLS data plane. [RFC9603] extends [RFC8664] to support SR for the IPv6 data plane.

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] section 5.1 specified 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, a PCE MUST NOT send SRv6 paths that exceed the SRv6 MSD capabilities of the PCC as specified in [RFC9603] section 5.1.

However, not all 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.

Besides, when a reduced SRH[RFC8754] is used, a reduced SRH does not contain the first segment of the related SR Policy, so the limitation of MSD could be loosen to allow one more SID in the SID list that is sent by the PCE.

This document analyzes the impact of different SRv6 MSDs when PCE sends SR-TE paths to the PCE. And it updates the MSD restrictions in [RFC9603] based on the analysis. This document also specifies a procedure for optimizing the number of SIDs in an SRv6-TE path that PCE can compute when the reduced SRH mode is used.

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. 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 reduced SRH is used.
- * 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, i.e, the segments Left is set to n-1 at the headend node(n is the number of elements in the SR Policy) regardless of reduced SRH is used.
- * 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 a "MUST" for PCE when sending the SRv6 path to PCC/ head-end node. And the limitation of Maximum H.Encaps MSD could be loosen when reduced SRH is used.

4. Updates to RFC9603

section 5.1 of [RFC9603] specified that "A PCE MUST NOT send SRv6 paths that exceed the SRv6 MSD capabilities of the PCC."

Based on the analysis in section 3. This document updates section 5.1 of RFC9603 to state that:

A PCE MUST NOT send SRv6 paths that exceed the SRv6 Maximum H.Encaps MSD and Maximum Segments Left MSD capabilities of the PCC.

5. New 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 TBD1):

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

6. Operational Considerations

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.

7. Security Considerations

Security considerations in [RFC9603] apply to this document.

8. IANA Considerations

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

Bit	Description	Reference

TBD1	Reduced SRH for MSD consideration	[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>>.
- [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

- [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>>.
- [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>>.
- [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>>.
- [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>>.

Author's Address

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