

Spring Working Group
Internet-Draft
Intended status: Standards Track
Expires: 29 May 2026

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

SRv6 Policy SID List Optimization
draft-ietf-spring-srv6-policy-sid-list-opt-00

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.

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 procedures 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/>.

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 29 May 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. Terminology	2
2. Introduction	3
3. SID List Optimization During Policy Creation	4
3.1. Policy Creation when using PCEP	4
3.2. Policy Creation when using BGP	4
4. SID List Optimization Advertisement	5
5. OAM Considerations	5
6. Operational Considerations	6
7. Security Considerations	6
8. IANA Considerations	6
9. Contributors	6
10. References	6
10.1. Normative References	7
10.2. Informative References	8
Authors' Addresses	9

1. Terminology

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

SR: Segment Routing.

SID: Segment Identifier.

SRv6: Segment Routing over IPv6 data plane.

2. 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. [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.

The SRv6 Policy 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.ietf-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 procedures needed to include or exclude the node SID when installing the SRv6 Policy.

3. SID List Optimization During Policy Creation

3.1. Policy Creation when using PCEP

The following procedure is applicable to both PCC-initiated Mode and PCE-initiated Mode.

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

PCC and PCE exchange capabilities during the PCEP initialization phase to indicate the support for the SID list optimization.

If the PCEP peers are capable of supporting the SID list optimization, the PCE indicates the inclusion or exclusion of the last SID in the ERO as follows:

- * If the computed SID list ends with the policy endpoint's Node SID and the traffic being steered over policy does not take the traffic to the policy endpoint, the PCE MUST signal the installation of the last node SID to the PCC. In this case, the PCC MUST install the last SID (the endpoint node SID) when installing the SRv6 Policy sid list(s) used to carry data traffic. In this case, the PCE considers the the policy endpoint's Node SID in the MSD consideration procedure. However, in the case when the installation of the last SID is not required, the PCE does not consider the the policy endpoint's Node SID in the MSD consideration procedure. Specifically, suppose the size of the SRv6 TE Policy SID list computed by PCE is L. In the case when the the last SID is not installed by PCC, the PCE uses the sid-list length L-1 in the headend MSD consideration procedure. This is because the endpoint node SID is suppressed.
- * If the computed SID list does not ends with the policy endpoint's Node SID, e.g., penultimate hop Adjacency SID, the PCC MUST install the last SID when installing the SRv6 Policy sid list(s) used to carry data traffic.

Protocol extension details can be found in [I-D.all-pce-srv6-policy-sid-list-optimization].

3.2. Policy Creation when using BGP

As defined in [RFC9256], an SR Policy is associated with one or more candidate paths. An SR Policy Controller [RFC9256] defines the set of policies and advertises them to policy headend routers (typically ingress routers). These policy advertisements use the BGP SR Policy SAFI [I-D.ietf-idr-sr-policy-safi].

When the SR Policy Controller completes computation of the SID list path:

- * If the computed SID list ends with the policy endpoint's Node SID and the traffic being steered over policy does not take the traffic to the policy endpoint, the SR Policy Controller installation of the last node SID to SRv6 Policy headend routers. In this case, the headend router MUST install the last SID (the endpoint node SID) when installing the SRv6 Policy sid list(s) used to carry data traffic.
- * If the computed SID list does not ends with the policy endpoint's Node SID, e.g., penultimate hop Adjacency SID, the SR Policy headend router MUST install this SID when installing the SRv6 Policy SID list(s) for data traffic.

Protocol extension details can be found in [I-D.liu-idr-sr-segment-list-optimize] and [I-D.lin-idr-sr-policy-admin-flags].

4. SID List Optimization Advertisement

[I-D.ietf-idr-bgp-ls-sr-policy] describes a mechanism to distribute SR policy information to external components using BGP-LS. SR policy information can be used by external components for path computation, re-optimization, service placement, network visualization, etc.

BGP-LS needs to be extended to indicate whether the endpoint's node SID is included or excluded in installing SID list(s) of the Candidate Path (CP) of an SRv6 Policy.

Protocol extension details can be found in [I-D.ali-idr-srv6-policy-sl-opt-distribution].

5. OAM Considerations

The SID list optimization outlined in this draft applies to the case where service traffic, like L3VPN, L2VPN, etc. is steered over policy. It does not apply to the probe packets used for OAM (Operational and Management) of the SRv6 policy in the transport network. When performing OAM operation over the policy, e.g., ping, traceroute, BFD, STAMP, etc. over policy, the initiator always includes the last SID of the policy segment list.

Please note that both data traffic and OAM packets are routed by the locator of the policy endpoint node. Therefore, the data traffic and OAM packets always follow the same route in the network.

6. Operational Considerations

This document specifies the detailed construction of the SRv6 Policy with inclusion or exclusion of end-point node SID and its operational mechanisms. Therefore, the manageability considerations of [RFC9256] apply.

Support of the inclusion or exclusion of end-point node SID capability defined in this document SHOULD be configurable. The control of this capability for the candidate path of an SRv6 Policy can be achieved via BGP, PCEP, YANG, or CLI. This capability itself does not alter packet forwarding behavior but rather optimizes the segment list orchestration for service packets, thereby enhancing forwarding performance.

A YANG model for the configuration and operation of SR Policy has been defined in [SR-POLICY-YANG]. In later versions, the capability to control the inclusion or exclusion of end-point node SID will be added under the candidate-path node of SRv6 Policy.

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

None

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

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

- [I-D.ali-idr-srv6-policy-sl-opt-distribution]
Ali, Z., Lin, C., Liu, Y., and R. Chen, "SRv6 Policy SID List Optimization Advertisement", Work in Progress, Internet-Draft, draft-ali-idr-srv6-policy-sl-opt-distribution-01, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ali-idr-srv6-policy-sl-opt-distribution-01>>.
- [I-D.all-pce-srv6-policy-sid-list-optimization]
Ali, Z., Lin, C., Liu, Y., Chen, R., and C. Li, "Path Computation Element Communication Protocol (PCEP) extensions for SRv6 Policy SID List Optimization", Work in Progress, Internet-Draft, draft-all-pce-srv6-policy-sid-list-optimization-03, 5 November 2025, <<https://datatracker.ietf.org/doc/html/draft-all-pce-srv6-policy-sid-list-optimization-03>>.
- [I-D.ietf-idr-bgp-ls-sr-policy]
Previdi, S., Talaulikar, K., Dong, J., Gredler, H., and J. Tantsura, "Advertisement of Segment Routing Policies using BGP Link-State", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ls-sr-policy-17, 6 March 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ls-sr-policy-17>>.
- [I-D.ietf-idr-sr-policy-safi]
Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and D. Jain, "Advertising Segment Routing Policies in BGP", Work in Progress, Internet-Draft, draft-ietf-idr-sr-policy-safi-13, 6 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-sr-policy-safi-13>>.
- [I-D.ietf-spring-srv6-mpls-interworking]
Agrawal, S., Filsfils, C., Voyer, D., Dawra, G., Li, Z., and S. Hegde, "SRv6 and MPLS interworking", Work in Progress, Internet-Draft, draft-ietf-spring-srv6-mpls-interworking-01, 7 July 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-mpls-interworking-01>>.

[I-D.lin-idr-sr-policy-admin-flags]

Lin, C., Liu, Y., Chen, R., and A. Karboubi, "BGP SR Policy Extensions for Administrative Flags", Work in Progress, Internet-Draft, draft-lin-idr-sr-policy-admin-flags-03, 5 November 2025, <<https://datatracker.ietf.org/doc/html/draft-lin-idr-sr-policy-admin-flags-03>>.

[I-D.liu-idr-sr-segment-list-optimize]

Liu, Y., Lin, C., and R. Chen, "BGP Extension for SRv6 Policy Segment List optimization", Work in Progress, Internet-Draft, draft-liu-idr-sr-segment-list-optimize-01, 20 February 2025, <<https://datatracker.ietf.org/doc/html/draft-liu-idr-sr-segment-list-optimize-01>>.

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