

IDR 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
7 July 2025

SRv6 Policy SID List Optimization Advertisement
draft-ali-idr-srv6-policy-sl-opt-distribution-01

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 BGP-LS extension 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.

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 8 January 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. Requirements Language	3
3. Terminology	3
4. Overview of BGP Extensions	4
5. Security Considerations	4
6. IANA Considerations	4
7. Contributors	4
8. References	4
8.1. Normative References	5
8.2. Informative References	5
Appendix A. Acknowledgements	5
Authors' Addresses	6

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). [I.D.draft-ietf-idr-bgp-ls-sr-policy] defines a mechanism to collect the Segment Routing Policy information that is locally available in a node and advertise it into BGP Link-State (BGP-LS) updates.

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-mppls-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 BGP-LS extension 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.

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

Headend node: Packet flows are steered into an SR Policy on a node where it is instantiated called a headend node [RFC9256].

SR: Segment Routing.

SID: Segment Identifier.

SRv6: Segment Routing over IPv6 data plane.

4. Overview of BGP Extensions

ENS (Endpoint Node Suppression flag) in the SR Candidate Path State TLV specified in [I-D- draft-ietf-idr-bgp-ls-sr-policy] is proposed to indicate whether the endpoint node SID is included or excluded in installing SID list(s) of the Candidate Path (CP). The flag is applicable only to SR policies with SRv6 data plane. The flag MUST NOT be set and MUST be ignored for SR policies with SR-MPLS data plane.

ENS (Endpoint Node Suppression flag) - 1 bit (Bit Position TBD1):

- * If set to 1 indicates the endpoint node SID is excluded when installing the SRv6 Policy SID list(s) used to carry the data traffic.
- * If set to 0 indicates the endpoint node SID is included when installing the SRv6 Policy SID list(s) used to carry the data traffic.

5. Security Considerations

TBA

6. IANA Considerations

TBA

7. 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

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

8. References

8.1. 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>>.
- [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>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <<https://www.rfc-editor.org/info/rfc9012>>.
- [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>>.

8.2. Informative References

- [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-00, 17 October 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-spring-srv6-mpls-interworking-00>>.

Appendix A. Acknowledgements

The authors would like to thank Ketan Talaulikar 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