

Interdomain Routing Working Group
Internet-Draft
Intended status: Standards Track
Expires: 23 April 2026

L. Gong
China Mobile
J. Dong
Huawei Technologies
C. Lin
New H3C Technologies
R. Chen
ZTE Corporation
20 October 2025

BGP Extensions for Inter-Domain Flexible Algorithm with Segment Routing
over IPv6 (SRv6)
draft-gong-idr-inter-domain-srv6-flex-algo-00

Abstract

IGP Flexible Algorithm (Flex-Algo) provides a mechanism for IGP nodes to compute the best paths according to a set of constraints in both the topology and computation metrics. Segment Routing over IPv6 (SRv6) can be used as one of the data plane of IGP Flex-Algo.

In SRv6 networks, services may be carried across multiple network domains which are under the same administration. For some services, it is expected that the an end-to-end inter-domain path can be computed according to the same constraints (such as low latency) defined by Flex-Algo.

This document describes BGP extensions to advertise SRv6 locators as IPv6 prefixes, together with the associated algorithm across the AS boundary.

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 23 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. The Scenario	3
3. Flex-Algo Extended Community	4
4. IANA Considerations	4
5. Security Considerations	4
6. Acknowledgements	5
7. References	5
7.1. Normative References	5
7.2. Informative References	5
Authors' Addresses	5

1. Introduction

IGP Flexible Algorithm (Flex-Algo) [RFC9350] provides a mechanism for IGP nodes to compute the best paths according to a set of constraints in the topology and computation metrics. Segment Routing over IPv6 (SRv6) [RFC8986] can be used as one data plane of IGP Flex-Algo.

In SRv6 networks, services may be carried across multiple network domains which are under the same administration. For some services, it is expected that the an end-to-end inter-domain path can be computed according to the same constraints (such as low latency) defined by Flex-Algo.

This document defines the BGP extensions to advertise SRv6 locators as IPv6 prefixes, together with the associated algorithm across the AS boundary.

2. The Scenario

This section shows a typical scenario of the inter-domain algorithm-specific IPv6 route distribution in Figure 1. AS1 and AS2 belong to the same network operator, the service from CE1 to CE2 requires an inter-domain path from PE1 to PE2 based on specific constraints (such as low latency). In each domain, the intra-domain low latency path can be computed based on flexible algorithm. In both AS1 and AS2, Flex-Algo 128 is defined to compute the best path using the latency metrics. In AS2, the SRv6 locator of PE2 which is associated with Flex-Algo 128 is advertised in IGP. To make the SRv6 locator of PE2 reachable in AS1, the SRv6 locator of PE2 is advertised as an IPv6 prefix by ASBR2 to ASBR1 in BGP IPv6 Unicast [RFC2545]. However, the associated algorithm is not carried in the Update message. When the BGP route is redistributed into IGP in AS1, the SRv6 locator of PE2 will be treated as a IPv6 prefix without algorithm information. As a result, the IGP nodes in AS1 will only compute the shortest path to the SRv6 locator of PE2 using the default IGP metric, then inter-domain path computed based on Flex-Algo 128 is not possible.





Figure 1. Inter-domain Flex-Algo specific route distribution

One possible approach to achieve inter-domain algorithm-specific path computation is to configure route-policy on ASBR1 to apply Flex-Algo 128 to specific prefixes received from ASBR2. This requires manual configuration of route-policy for each prefix which is algorithm-specific SRv6 locators, thus it is complex and error-prone. A preferable approach is to add the associated algorithm information in the BGP routes advertised across the AS boundary, so that when the BGP route is redistributed into IGP, the associated algorithm could be obtained automatically.

3. Flex-Algo Extended Community

A new extended community is defined to carry the algorithm information associated with the prefix carried in the IPv6 Unicast NLRI. The Flex-Algo extended community is a transitive Opaque Extended Community which is used to carry the Algorithm ID associated with the IPv6 prefixes advertised in BGP Update message. The encoding of the Flex-Algo Extended Community is shown as below:



The value of the extended Type field is 0x03, which indicates it is a transitive opaque extended community. The value of the low-order octet of the extended Type field is TBA.

Flags: 1 octet, carries the flags related to an Algorithm. No flag is defined in this document.

Algorithm: 1 octet, it can be the algorithm as defined in the "IGP Algorithm Types" registry [RFC8665], or Flexible Algorithms as defined in [RFC9350].

Reserved: 4-octet reserved for future use.

4. IANA Considerations

IANA is requested to assign a code point for Flex-Algo Extended Community from the "Transitive Opaque Extended Community Sub-Types" subregistry of the "Border Gateway Protocol (BGP) Extended Communities" registry:

Sub-type Value	Name	Reference
TBA	Flex-Algo Extended Community	This document

5. Security Considerations

This document does not introduce any additional security considerations than those described in [RFC4271] and [RFC4272].

6. Acknowledgements

The authors would like to thank XXX for the review and valuable discussion.

7. References

7.1. Normative References

- [RFC2545] Marques, P. and F. Dupont, "Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing", RFC 2545, DOI 10.17487/RFC2545, March 1999, <<https://www.rfc-editor.org/info/rfc2545>>.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <<https://www.rfc-editor.org/info/rfc4272>>.
- [RFC8665] Psenak, P., Ed., Previdi, S., Ed., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", RFC 8665, DOI 10.17487/RFC8665, December 2019, <<https://www.rfc-editor.org/info/rfc8665>>.
- [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>>.
- [RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <<https://www.rfc-editor.org/info/rfc9350>>.

7.2. Informative References

Authors' Addresses

Liyan Gong
China Mobile
China
Email: gongliyan@chinamobile.com

Jie Dong
Huawei Technologies
China
Email: jie.dong@huawei.com

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

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