

IDR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 6 May 2026

Y. Liu
Shaofu. Peng
ZTE Corporation
2 November 2025

Advertisement of Algorithm in BGP
draft-lp-idr-bgp-algorithm-00

Abstract

This document proposes extensions to BGP to support algorithm-based end-to-end path establishment.

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 6 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. Introduction	2
2. Conventions used in this document	2
2.1. Requirements Language	2
3. Algorithm Extended Community	2
4. Algo-based Inter-domain Path	3
5. Security Considerations	4
6. IANA Considerations	5
7. References	5
7.1. Normative References	5
7.2. Informative References	5
Authors' Addresses	6

1. Introduction

[RFC9350] proposes a solution that allows IGPs themselves to compute constraint-based paths over the SR network. [RFC9502] allows flex-algo to be deployed in any IP network, even in the absence of SR-MPLS and SRv6.

However, the algorithm-based path can only be used in the IGP domain. In the BGP-based inter-domain scenario, end-to-end path based on algorithms cannot be supported.

This document proposes extensions to BGP to support algorithm-based end-to-end path establishment.

2. Conventions used in this document

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. Algorithm Extended Community

This document defines a new transitive BGP Extended Communities Attribute[RFC4360]. This new Extended Community has the following encoding, where:




Figure 1: Algorithm Extended Community

Type: 1 octet. The value is 0x03.

Sub-Type: 1 octet. TBA1

Flags: 1 octet. Unused, MUST be set to 0 and ingored on receipt.

Algorithm: 1 octet specifying IGP Algorithm Types. Value from 0 to 255.

Reserved: 4 Octets. MUST be set to 0, ignored at reception

4. Algo-based Inter-domain Path




Figure 2: Algo-based Inter-domain Path

As shown in Figure 2, node 8 is configured with two loopback addresses, loopback-1 and loopback-2, they belong to the flex-algo 128 plane and the flex-algo 129 plane respectively. In IGP domain 2, the routes to loopback-1 will be generated on the nodes (e.g, node 5,6,8) in the Flex-algo 128 plane and the routes to loopback-2 will be generated on the nodes (e.g, node 5,7,8) in the Flex-algo 129 plane.

Node 5 can advertise prefix loopback-1 and prefix loopback-2 to node 4 through BGP[RFC2545][RFC4271]. Node 4 can import the BGP routes into IGP and continue to advertise the routes to its neighbors in IGP domain1. Or, node 4 directly advertises the routes to node 1 through BGP.

In both cases, the corresponding algorithm information from IGP domain2 is lost during the advertisement. As a result, node 4 does not know which Flex-algo plane to import loopback-1 or loopback-2 into IGP domain 1.

With the Algorithm Extended Community, the algorithm information can be carried in the BGP route of loopbacks advertised from node 5 to node 4.

The administrator can configure algorithms in each IGP domain in the network. A simple configuration method is that algorithms in each IGP domain are consistent. If the algorithm configurations in each IGP domain are inconsistent, the ASBR needs to know the mapping relationship of the algorithms and carry the converted algorithm information in Algorithm Extended Community when advertising the BGP route.

A BGP speaker can advertise multiple paths for the same address prefix, each path is identified by a Path Identifier in addition to the address prefix [RFC7911]. By leveraging add-path, multiple loopbacks on the egress node can be avoided.

Same approach is applicable for SRv6 locator which is also advertised via [RFC2545], SR-MPLS BGP Prefix SID advertisement[RFC8669] and BGP Labeled Unicast(BGP-LU)[RFC8227]. If any Router Reflector existed in the network, it SHOULD support this new Extended Community.

5. Security Considerations

TBD

6. IANA Considerations

IANA is requested to allocate the sub-type TBA1 for "Algorithm Extended Community" under the "BGP Transitive Opaque Extended Community"

Sub-type Value	Name	Reference
TBA1	Algorithm Extended Community	This document

7. References

7.1. 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>>.
- [RFC4360] Sangli, S., Tappan, D., and Y. Rekhter, "BGP Extended Communities Attribute", RFC 4360, DOI 10.17487/RFC4360, February 2006, <<https://www.rfc-editor.org/info/rfc4360>>.
- [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <<https://www.rfc-editor.org/info/rfc7911>>.
- [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>>.
- [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>>.
- [RFC9502] Britto, W., Hegde, S., Kaneriya, P., Shetty, R., Bonica, R., and P. Psenak, "IGP Flexible Algorithm in IP Networks", RFC 9502, DOI 10.17487/RFC9502, November 2023, <<https://www.rfc-editor.org/info/rfc9502>>.

7.2. Informative 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>>.
- [RFC8227] Cheng, W., Wang, L., Li, H., van Helvoort, H., and J. Dong, "MPLS-TP Shared-Ring Protection (MSRP) Mechanism for Ring Topology", RFC 8227, DOI 10.17487/RFC8227, August 2017, <<https://www.rfc-editor.org/info/rfc8227>>.
- [RFC8669] Previdi, S., Filsfils, C., Lindem, A., Ed., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix Segment Identifier Extensions for BGP", RFC 8669, DOI 10.17487/RFC8669, December 2019, <<https://www.rfc-editor.org/info/rfc8669>>.

Authors' Addresses

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

Peng Shaofu
ZTE Corporation
Email: peng.shaofu@zte.com.cn