

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

L. Ginsberg
P. Psenak
Cisco Systems
Z. Zhang
ZTE Corporation
20 October 2025

IGP Flex Soft Dataplane
draft-ietf-lsr-flex-soft-dataplane-00

Abstract

Advertisement of IGP Flex-Algo participation requires a dataplane context. This document defines a "soft dataplane" usable in cases where existing defined dataplanes are not suitable.

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. Requirements Language	2
3. Use Case Discussion	2
4. IS-IS Soft Dataplane Algorithm Sub-TLV	3
5. OSPF Soft Dataplane Algorithm TLV	4
6. IANA Considerations	5
6.1. OSPF Router Information TLV Registry	5
6.2. IS-IS Router Capability sub-TLV Registry	6
6.3. Security Considerations	6
7. Normative References	6
Authors' Addresses	7

1. Introduction

Advertisement of IGP Flex-Algo[RFC9350] participation requires a dataplane context. Existing data planes which have been defined include:

Segment Routing Dataplane [RFC8667] [RFC8665]

IP Flex Dataplane [RFC9502]

The need to use an IGP Flexible Algorithm may occur in deployments where none of the existing dataplanes are supported or suitable.

In such cases a "soft dataplane" MAY be used to provide the necessary context for advertisement of Flex-Algo support. This document defines the mechanisms to advertise such a dataplane.

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. Use Case Discussion

A deployment may require the use of flex-algo to achieve traffic flows that meet certain constraints. In some cases, flex-algo paths may be used by an application that does not require the use of any of the currently defined dataplanes supported by flex-algo. The use of these dataplanes may not be desired and/or is not supported in the network. IP Flex-algo extends flex-algo so that it can be used directly with IPv4 and IPv6 forwarding, but the use of IPv4/IPv6

Algorithm specific prefix advertisements may not always be possible, especially if the existing deployment does not allow for the lack of reachability using the base IGP SPF (AKA algorithm 0) for these prefixes. The use of a soft dataplane provides context which still allows flex-algo to be deployed in such cases.

It might be thought that the use of multiple topologies [RFC5120] [RFC4915] would suffice for such use cases, but such a viewpoint overlooks the fundamental capabilities that flex-algo provides. These include the use of a calculation type, metric type, and/or link attributes that either differ from what is natively supported by the IGP or are topology independent attributes.

For example, one use case might be to support multicast distribution over a constrained topology in an IP only network. The use case may want to provide a low latency multicast service, in which case the multicast tree needs to be built using latency as a metric. Flex-algo can provide the best path calculation that is needed for such a use case.

The new dataplane is referred to as "soft" because the flex-algo paths computed for this dataplane are not expected to be used by forwarding directly - e.g., they will not be installed in the data path. They may be used by an application to create a forwarding state that is maintained by the application itself.

Note that multiple flex algorithms - each defined for a different use case - can be advertised in the context of a single soft dataplane. Therefore, it is expected that a single soft dataplane will suffice for all possible use cases.

The following sections define how to advertise flex-algorithm support in the context of the soft dataplane.

4. IS-IS Soft Dataplane Algorithm Sub-TLV

The IS-IS [ISO10589] Soft Dataplane Algorithm Sub-TLV is a sub-TLV of the IS-IS Router Capability TLV [RFC7981] and has the following format:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
      +-----+-----+-----+-----+
      |  Type   |          Length          |
      +-----+-----+-----+-----+
      | Algorithm 1 | Algorithm 2 | Algorithm ... | Algorithm n |
      +-----+-----+-----+-----+

```

Figure 1: IS-IS Soft Dataplane Algorithm Sub-TLV

Type (1 octet): Soft Dataplane Algorithm Sub-TLV (Value TBD)

Length (1 octet): Variable

Algorithm (1 octet): Value from 128 to 255

The IS-IS Soft Dataplane Algorithm Sub-TLV MUST be propagated throughout the level and MUST NOT be advertised across level boundaries. Therefore, the S bit in the Router Capability TLV, in which the IS-IS Soft Dataplane Algorithm Sub-TLV is advertised, MUST NOT be set.

The IS-IS Soft Dataplane Algorithm Sub-TLV is optional. It MUST NOT be advertised more than once at a given level. A router receiving multiple IS-IS Soft Dataplane Algorithm sub-TLVs from the same originator MUST select the first advertisement in the lowest-numbered Link State PDU (LSP), and subsequent instances of the IS-IS Soft Dataplane Algorithm Sub-TLV MUST be ignored.

Algorithms outside the Flex-Algorithm range (128-255) MUST be ignored by the receiver. This situation SHOULD be logged as an error.

The Flex-Algorithm participation advertised in the IS-IS Soft Dataplane Algorithm Sub-TLV is topology independent. When a router advertises participation in the IS-IS Soft Dataplane Algorithm Sub-TLV, the participation applies to all topologies in which the advertising node participates.

5. OSPF Soft Dataplane Algorithm TLV

The OSPF [RFC2328] Soft Dataplane Algorithm TLV is a top-level TLV of the Router Information Opaque Link State Advertisement (LSA) [RFC7770] and has the following format:




Figure 2: OSPF Soft Dataplane Algorithm TLV

Type (2 octets): Soft Dataplane Algorithm TLV (TBD)

Length(2 octets): Variable

Algorithm (1 octet): Value from 128 to 255

The OSPF Soft Dataplane Algorithm TLV is optional. It MUST only be advertised once in the Router Information LSA.

Algorithms outside the Flex-Algorithm range (128-255) MUST be ignored by the receiver. This situation SHOULD be logged as an error.

When multiple OSPF Soft Dataplane Algorithm TLVs are received from a given router, the receiver MUST use the first occurrence of the TLV in the Router Information LSA. If the OSPF Soft Dataplane Algorithm TLV appears in multiple Router Information LSAs that have different flooding scopes, the OSPF Soft Dataplane Algorithm TLV in the Router Information LSA with the area-scoped flooding scope MUST be used. If the OSPF Soft Dataplane Algorithm TLV appears in multiple Router Information LSAs that have the same flooding scope, the OSPF Soft Dataplane Algorithm TLV in the Router Information LSA with the numerically smallest Instance ID (Opaque ID for OSPFv2 or Link State ID for OSPFv3) MUST be used, and subsequent instances of the OSPF Soft Dataplane Algorithm TLV MUST be ignored.

The Router Information LSA can be advertised at any of the defined flooding scopes (link, area, or Autonomous System (AS)). For the purpose of OSPF Soft Dataplane Algorithm TLV advertisement, area or AS-scoped flooding is REQUIRED. The AS flooding scope SHOULD NOT be used unless local configuration policy on the originating router indicates domain-wide flooding.

The Flexible Algorithm participation advertised in the OSPF Soft Dataplane Algorithm TLV is topology independent. When a router advertises participation in an OSPF Soft Dataplane Algorithm TLV, the participation applies to all topologies in which the advertising node participates.

6. IANA Considerations

6.1. OSPF Router Information TLV Registry

This document updates the "OSPF Router Information (RI) TLVs" registry as follows:

Value	TLV Name	Reference
TBD	Soft Dataplane Algorithm	draft-ginsberg-lsr-soft-dataplane

Table 1

6.2. IS-IS Router Capability sub-TLV Registry

This document updates the "IS-IS Sub-TLVs for IS-IS Router CAPABILITY TLV" registry as follows:

Value	TLV Name	Reference
TBD	Soft Dataplane Algorithm	draft-ginsberg-lsr-soft-dataplane

Table 2

6.3. Security Considerations

This document creates no new security issues for the IGPs.

7. Normative References

- [ISO10589] ISO, "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", August 1987, <ISO/IEC 10589:2002>.
- [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>>.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007, <<https://www.rfc-editor.org/info/rfc4915>>.

- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.
- [RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 7770, DOI 10.17487/RFC7770, February 2016, <<https://www.rfc-editor.org/info/rfc7770>>.
- [RFC7981] Ginsberg, L., Previdi, S., and M. Chen, "IS-IS Extensions for Advertising Router Information", RFC 7981, DOI 10.17487/RFC7981, October 2016, <<https://www.rfc-editor.org/info/rfc7981>>.
- [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>>.
- [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>>.
- [RFC8667] Previdi, S., Ed., Ginsberg, L., Ed., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", RFC 8667, DOI 10.17487/RFC8667, December 2019, <<https://www.rfc-editor.org/info/rfc8667>>.
- [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., Kaneriy, 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>>.

Authors' Addresses

Les Ginsberg
Cisco Systems
Email: ginsberg@cisco.com

Peter Psenak
Cisco Systems
Apollo Business Center
Mlynske nivy 43
82109 Bratislava
Slovakia
Email: ppsenak@cisco.com

Zheng Zhang
ZTE Corporation
China
Email: zhang.zheng@zte.com.cn