LSR Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 17, 2022

W. Britto
S. Hegde
P. Kaneriya
R. Shetty
R. Bonica
Juniper Networks
P. Psenak
Cisco Systems
May 16, 2022

IGP Flexible Algorithms (Flex-Algorithm) In IP Networks
draft-ietf-lsr-ip-flexalgo-06

Abstract

An IGP Flexible Algorithm (Flex-Algorithm) allows IGPs to compute constraint-based paths. The base IGP Flex-Algorithm specification describes how it is used with Segment Routing (SR) data planes - SR MPLS and SRv6.

This document extends IGP Flex-Algorithm, so that it can be used with regular IPv4 and IPv6 forwarding.

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 November 17, 2022.

Table of Contents

1. Introduction

An IGP Flex-Algorithm as specified in [I-D.ietf-lsr-flex-algo] computes a constraint-based path to:

- All Flex-Algorithm specific Prefix Segment Identifiers (SIDs) [RFC8402].

- All Flex-Algorithm specific SRv6 Locators [RFC8986].

Therefore, Flex-Algorithm cannot be deployed in the absence of SR and SRv6.

This document extends Flex-Algorithm, allowing it to compute paths to IPv4 and IPv6 prefixes.

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 Example

Mobile networks are becoming more and more IP centric. Each end-user session from a gNB (gNodeB) can be destined to a specific UPFs (User Plane Function) based on the session requirements. For example, some sessions require high bandwidth, others need to be routed along the lowest latency path. Each UPF is assigned a unique IP address. As a result, traffic for different sessions is destined to a different destination IP address.

By associating the prefix that contains the UPF addresses with a specific IP algorithm and routing the algorithm specific traffic according to a certain constraints, e.g., low latency, a session traffic is routed according to the SLA (Service Level Agreement) appropriate for such session.

4. Advertising Flex-Algorithm Definitions (FAD)

To guarantee loop free forwarding, all routers that participate in a Flex-Algorithm MUST agree on the Flex-Algorithm Definition (FAD).

Selected nodes within the IGP domain MUST advertise FADs as described in Sections 5, 6, and 7 of [I-D.ietf-lsr-flex-algo].

5. Advertising IP Flex-Algorithm Participation

A node may use various algorithms when calculating paths to nodes and prefixes. Algorithm values are defined in the IGP Algorithm Type Registry [IANA-ALG].

A node MUST participate in a Flex-Algorithm to be:

- Able to compute path for such Flex-Algorithm
- Part of the topology for such Flex-Algorithm

Flex-Algorithm participation MUST be advertised for each Flex-Algorithm data-plane independently, as specified in [I-D.ietf-lsr-flex-algo]. Using Flex-Algorithm for regular IPv4 and IPv6 prefixes represents an independent Flex-Algorithm data-plane, and as such, the Flex-Algorithm participation for the IP Flex-Algorithm data-plane MUST be signalled independently of any other Flex-Algorithm data-plane (e.g., SR).

Advertisement of participation in IP Flex-Algorithm MUST NOT impact the router participation in default algorithm 0.

Advertisement of participation in IP Flex-Algorithm MUST NOT impact the router participation signaled for other data-planes.

The following sections describe how the IP Flex-Algorithm participation is advertised in IGP protocols.

5.1. The IS-IS IP Algorithm Sub-TLV

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

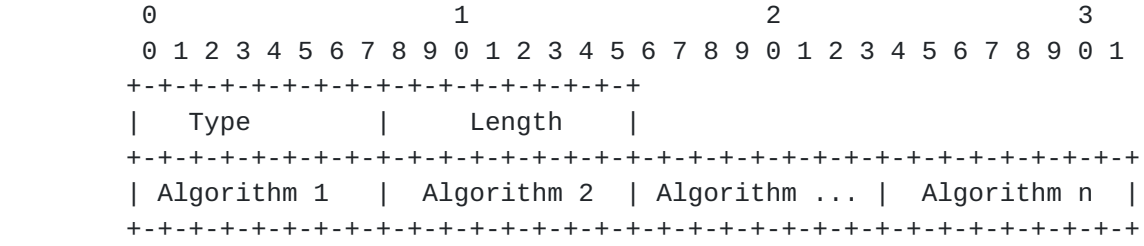

Figure 1: IS-IS IP Algorithm Sub-TLV

- Type: IP Algorithm Sub-TLV (Value 29)
- Length: Variable
- Algorithm (1 octet): Value from 128 to 255.

The IP 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 IP Algorithm Sub-TLV is advertised, MUST NOT be set.

The IP Algorithm Sub-TLV is optional. It MUST NOT be advertised more than once at a given level. A router receiving multiple IP Algorithm sub-TLVs from the same originator MUST select the first advertisement

in the lowest-numbered LSP and subsequent instances of the IP Algorithm Sub-TLV MUST be ignored.

The use of IP Algorithm Sub-TLV to advertise support for algorithms outside the Flex-Algorithm range (128-255) is outside the scope of this document.

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

5.2. The OSPF IP Algorithm TLV

The OSPF [RFC2328] IP Algorithm TLV is a top-level TLV of the Router Information Opaque LSA [RFC7770] and has the following format:

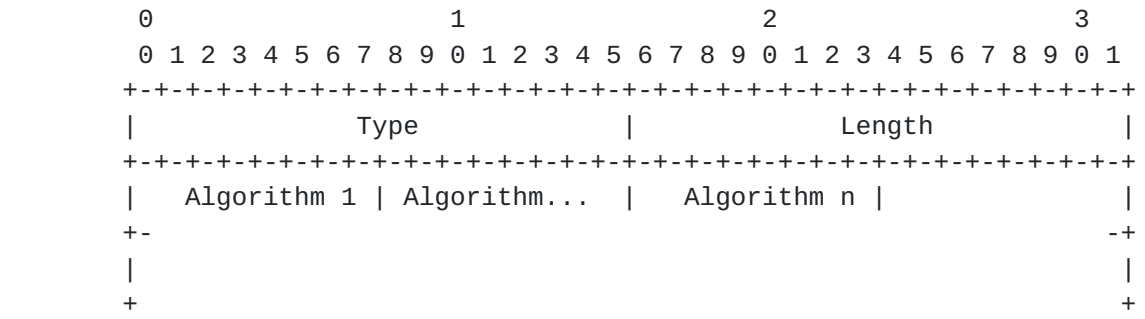

Figure 2: OSPF IP Algorithm TLV

- Type: IP Algorithm TLV (Value TBD by IANA)
- Length: Variable
- Algorithm (1 octet): Value from 128 to 255.

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

When multiple IP 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 IP Algorithm TLV appears in multiple Router Information LSAs that have different flooding scopes, the IP Algorithm TLV in the Router Information LSA with the area-scoped flooding scope MUST be used. If the IP Algorithm TLV appears in multiple Router Information LSAs that have the same flooding scope, the IP 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 IP 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 IP Algorithm TLV advertisement, area-scoped flooding is REQUIRED.

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

6. Advertising IP Flex-Algorthm Reachability

To be able to associate the prefix with the Flex-Algorithm, the existing prefix reachability advertisements can not be used, because they advertise the prefix reachability in default algorithm 0. Instead, a new IP Flex-Algorithm reachability advertisements are defined in IS-IS and OSPF.

The M-flag in FAD is not applicable to IP Algorithm Prefixes. Any IP Algorithm Prefix advertisement includes the Algorithm and Metric fields. When an IP Algorithm Prefix is advertised between areas or domains, the metric field in the IP Algorithm Prefix advertisement MUST be used irrespective of the M-flag in the FAD advertisement.

6.1. The IS-IS IPv4 Algorithm Prefix Reachability TLV

A top-level TLV is defined for advertising IPv4 Flex-Algorithm Prefix Reachability in IS-IS - IPv4 Algorithm Prefix Reachability TLV.

This new TLV shares the sub-TLV space defined for TLVs Advertising Prefix Reachability.

The IS-IS IPv4 Algorithm Prefix Reachability TLV has the following format:

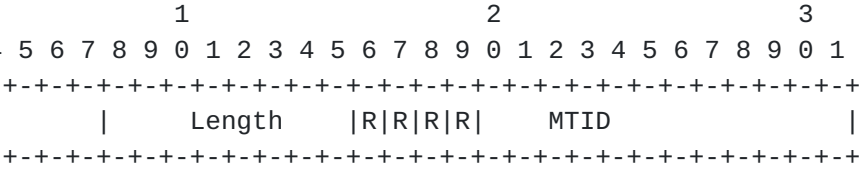

Figure 3: IS-IS IPv4 Algorithm Prefix Reachability TLV

o Type: IPv4 Algorithm Prefix Reachability TLV (Value 126).

- Length: Variable.

- R bits (4 bits): Reserved for future use. They MUST be set to zero on transmission and MUST be ignored on receipt.

- MTID (12 bits): Multitopology Identifier as defined in [RFC5120]. Note that the value 0 is legal.

Followed by one or more prefix entries of the form:

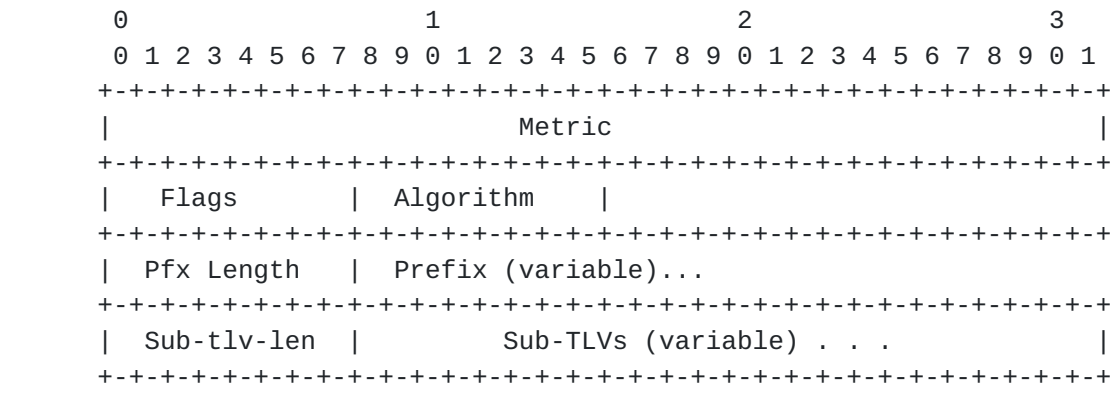

Figure 4: IS-IS IPv4 Algorithm Prefix Reachability TLV

- Metric (4 octets): Metric information.

- Flags (1 octet):

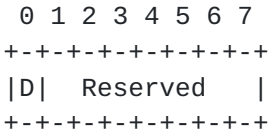

  D-flag: When the Prefix is leaked from level-2 to level-1, the D bit MUST be set. Otherwise, this bit MUST be clear. Prefixes with the D bit set MUST NOT be leaked from level-1 to level-2. This is to prevent looping.

- Algorithm (1 octet): Associated Algorithm from 128 to 255.

- Prefix Len (1 octet): Prefix length measured in bits.

- Prefix (variable length): Prefix mapped to Flex-Algorithm.

- Optional Sub-TLV-length (1 octet): Number of octets used by sub-TLVs

- Optional sub-TLVs (variable length).

A router receiving multiple IPv4 Algorithm Prefix Reachability advertisements for the same prefix, from the same originator, each with a different Algorithm, MUST select the first advertisement in the lowest-numbered LSP and ignore any subsequent IPv4 Algorithm Prefix Reachability advertisements for the same prefix for any other Algorithm.

A router receiving multiple IPv4 Algorithm Prefix Reachability advertisements for the same prefix, from different originators, each with a different Algorithm, MUST ignore all of them and MUST NOT install any forwarding entries based on these advertisements. This situation SHOULD be logged as an error.

In cases where a prefix advertisement is received in both a IPv4 Prefix Reachability TLV and an IPv4 Algorithm Prefix Reachability TLV, the IPv4 Prefix Reachability advertisement MUST be preferred when installing entries in the forwarding plane.

6.2. The IS-IS IPv6 Algorithm Prefix Reachability TLV

The IS-IS IPv6 Algorithm Prefix Reachability TLV is identical to the IS-IS IPv4 Algorithm Prefix Reachability TLV, except that it has a unique type. The type is 127.

A router receiving multiple IPv6 Algorithm Prefix Reachability advertisements for the same prefix, from the same originator, each with a different Algorithm, MUST select the first advertisement in the lowest-numbered LSP and ignore any subsequent IPv6 Algorithm Prefix Reachability advertisements for the same prefix for any other Algorithm.

A router receiving multiple IPv6 Algorithm Prefix Reachability advertisements for the same prefix, from different originators, each with a different Algorithm, MUST ignore all of them and MUST NOT install any forwarding entries based on these advertisements. This situation SHOULD be logged as an error.

In cases where a prefix advertisement is received in both an IPv6 Prefix Reachability TLV and an IPv6 Algorithm Prefix Reachability TLV, the IPv6 Prefix Reachability advertisement MUST be preferred when installing entries in the forwarding plane.

In cases where a prefix advertisement is received in both IS-IS SRv6 Locator TLV and in IS-IS IPv6 Algorithm Prefix Reachability TLV, the receiver MUST ignore both of them and MUST NOT install any forwarding entries based on these advertisements. This situation SHOULD be logged as an error.

6.3. The OSPFv2 IP Algorithm Prefix Reachability Sub-TLV

A new Sub-TLV of the OSPFv2 Extended Prefix TLV is defined for advertising IP Algorithm Prefix Reachability in OSPFv2, the OSPFv2 IP Algorithm Prefix Reachability Sub-TLV.

The OSPFv2 IP Algorithm Prefix Reachability Sub-TLV has the following format:

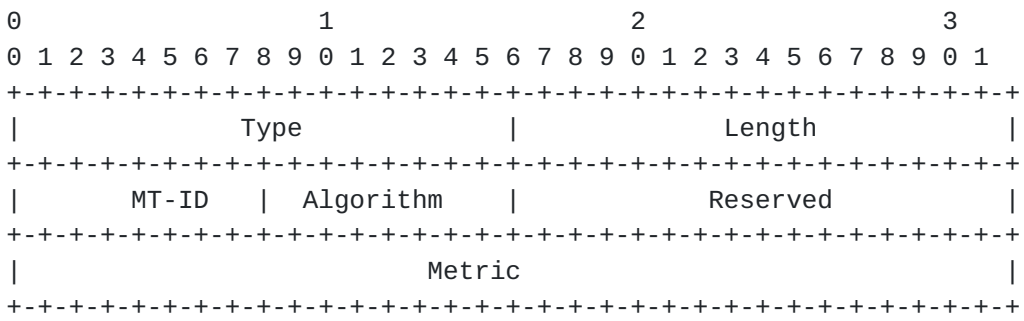

Figure 5: OSPFv2 IP Algorithm Prefix Reachability Sub-TLV

- Type (2 octets) : The value is TBD.
- Length (2 octet): 8
- MT-ID (1 octet): Multi-Topology ID as defined in [RFC4915]
- Algorithm (1 octet): Associated Algorithm from 128 to 255.
- Reserved: (2 octets). SHOULD be set to 0 on transmission and MUST be ignored on reception.
- Metric (4 octets): The algorithm specific metric value.

An OSPFv2 router receiving multiple OSPFv2 IP Algorithm Prefix Reachability Sub-TLVs in the same OSPFv2 Extended Prefix TLV, MUST select the first advertisement of this Sub-TLV and MUST ignore all remaining occurences of this Sub-TLV in the OSPFv2 Extended Prefix TLV.

An OSPFv2 router receiving multiple OSPFv2 IP Algorithm Prefix Reachability TLVs for the same prefix, from different originators, each with a different Algorithm, MUST ignore all of them and MUST NOT install any forwarding entries based on these advertisements. This situation SHOULD be logged as an error.

In cases where a prefix advertisement is received in any of the LSAs advertising the prefix reachability for algorithm 0 and in an OSPFv2 IP Algorithm Prefix Reachability TLV, only the prefix reachability advertisement for algorithm 0 MUST be used and all occurences of the OSPFv2 IP Algorithm Prefix Reachability TLV MUST be ignored.

6.4. The OSPFv3 IP Algorithm Prefix Reachability Sub-TLV

The OSPFv3 [RFC5340] IP Algorithm Prefix Reachability Sub-TLV is defined for advertisement of the IP Algorithm Prefix Reachability in OSPFv3.

The OSPFv3 IP Algorithm Prefix Reachability Sub-TLV is a sub-TLV of the following OSPFv3 TLVs defined in [RFC8362]:

- Intra-Area-Prefix TLV
- Inter-Area-Prefix TLV
- External-Prefix TLV

The format of OSPFv3 IP Algorithm Prefix Reachability Sub-TLV is shown below:

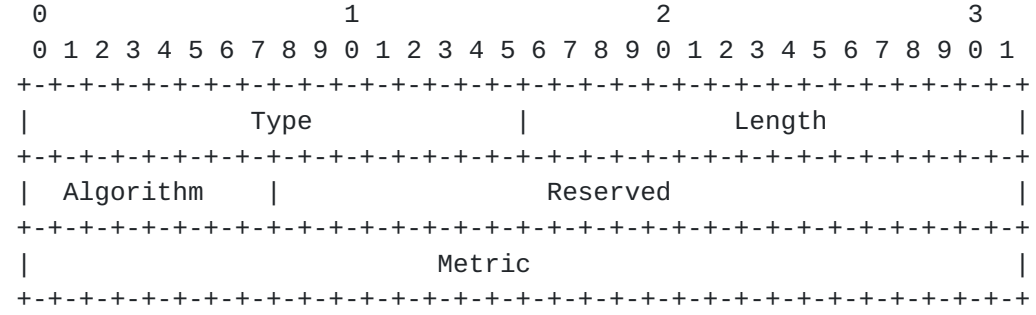

Figure 6: OSPFv3 IP Algorithm Prefix Reachability Sub-TLV

Where:

Type (2 octets): The value is TBD.

Length (2 octets): 8.

Algorithm (1 octet): Associated Algorithm from 128 to 255.

Reserved: (3 octets). SHOULD be set to 0 on transmission and MUST be ignored on reception.

Metric (4 octets): The algorithm specific metric value.

When the OSPFv3 IP Algorithm Prefix Reachability Sub-TLV is present, the metric value in its parent TLV MUST be set to LSInfinity [RFC2328]. If the metric value in the parent TLV is not set to LSInfinity, the OSPFv3 IP Algorithm Prefix Sub-TLV MUST be ignored by the receiver.

An OSPFv3 router receiving multiple OSPFv3 IP Algorithm Prefix Reachability Sub-TLVs in the same parent TLV, MUST select the first advertisement of this Sub-TLV and MUST ignore all remaining occurences of this Sub-TLV in the parent TLV.

An OSPFv3 router receiving multiple OSPFv3 IP Algorithm Prefix Reachability TLVs for the same prefix, from different originators, each with a different Algorithm, MUST ignore all of them and MUST NOT install any forwarding entries based on these advertisements. This situation SHOULD be logged as an error.

In cases where a prefix advertisement is received in any of the LSAs advertising the prefix reachability for algorithm 0 and in an OSPFv3 OSPFv3 IP Algorithm Prefix Reachability Sub-TLV, only the prefix reachability advertisement for algorithm 0 MUST be used and all occurences of the OSPFv3 IP Algorithm Prefix Reachability TLV MUST be ignored.

In cases where a prefix advertisement is received in both an OSPFv3 SRv6 Locator TLV and in an OSPFv3 IP Algorithm Prefix Reachability Sub-TLV, the receiver MUST ignore both of them and MUST NOT install any forwarding entries based on these advertisements. This situation SHOULD be logged as an error.

6.5. The OSPF IP Flexible Algorithm ASBR Metric Sub-TLV

[I-D.ietf-lsr-flex-algo] defines the OSPF Flexible Algorithm ASBR Metric Sub-TLV (FAAM) that is used by an OSPFv2 or an OSPFv3 ABR to advertise a Flex-Algorithm specific metric associated with the corresponding ASBR LSA.

As described in [I-D.ietf-lsr-flex-algo] each data-plane signals the participation independently. IP Flex-Algorithm participation is signaled independent of the Segment Routing (SR) Flex-Algorithm participation. As a result, the calculated topologies for SR and IP Flex-Algorithm could be different. Such difference prevents the usage of FAAM for the purpose of the IP Flex-Algorithm.

The OSPF IP Flexible Algorithm ASBR Metric (IPFAAM) Sub-TLV is defined for the advertisement of the IP Flex-Algorithm specific metric associated with an ASBR by the ABR.

The IPFAAM Sub-TLV is a Sub-TLV of the:

- OSPFv2 Extended Inter-Area ASBR TLV as defined in [I-D.ietf-lsr-flex-algo]

- OSPFv3 Inter-Area-Router TLV defined in [RFC8362]

The OSPF IPFAAM Sub-TLV has the following format:

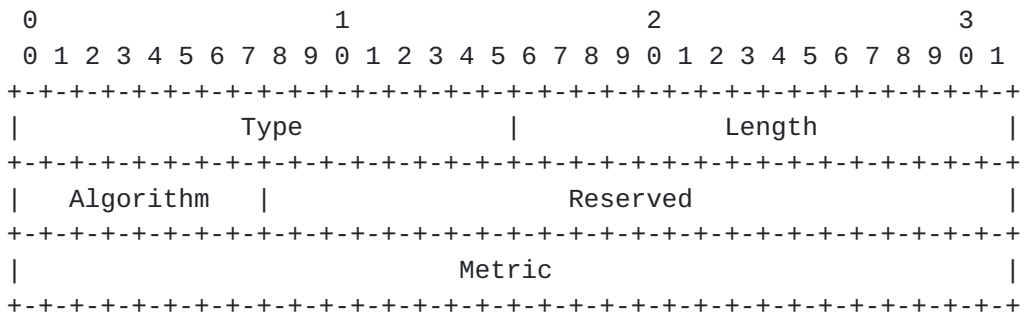

where:

Figure 7: OSPF IP Flexible Algorithm ASBR Metric Sub-TLV

Type (2 octets): TBD for OSPFv2, TBD for OSPFv3.

Length (2 octets): 8.

Algorithm (1 octet): Associated Algorithm from 128 to 255.

Reserved: (3 octets). SHOULD be set to 0 on transmission and MUST be ignored on reception.

Metric (4 octets): The algorithm specific metric value.

The usage of the IPFAAM Sub-TLV is similar to the usage of the FAAM Sub-TLV defined in [I-D.ietf-lsr-flex-algo], but it is used to advertise IP Flex-Algorithm metric.

An OSPF ABR MUST include the OSPF IPFAAM Sub-TLVs as part of the ASBR reachability advertisement between areas for every IP Flex-Algorithm in which it participates and the ASBR is reachable in.

The FAAM Sub-TLV as defined in [I-D.ietf-lsr-flex-algo] MUST NOT be used during IP Flex-Algorithm path calculation, the IPFAAM Sub-TLV MUST be used instead.

7. Calculating of IP Flex-Algorthm Paths

The IP Flex-Algorthm is considered as yet another data-plane of the Flex-Algorithm as described [I-D.ietf-lsr-flex-algo].

Participation for the IP Flex-Algorithm is signalled as described in Section 5 and is specific to the IP Flex-Algorithm data-plane.

Calculation of IP Flex-Algorithm paths follows what is described in [I-D.ietf-lsr-flex-algo]. This computation uses the IP Flex-Algorithm data-plane participation and is independent of the Flex-Algorithm calculation done for any other Flex-Algorithm data-plane (e.g., SR, SRv6).

The IP Flex-Algorithm data-plane only considers participating nodes during the Flex-Algorithm calculation. When computing paths for a given Flex-Algorithm, all nodes that do not advertise participation for the IP Flex-Algorithm, as described in Section 5, MUST be pruned from the topology.

8. IP Flex-Algorthm Forwarding

The IP Algorithm Prefix Reachability advertisement as described in Section 5 includes the MTID value that associates the prefix with a specific topology. Algorithm Prefix Reachability advertisement also includes an Algorithm value that explicitly associates the prefix with a specific Flex-Algorithm. The paths to the prefix MUST be calculated using the specified Flex-Algorithm in the associated topology.

Forwarding entries for the IP Flex-Algorithm prefixes advertised in IGPs MUST be installed in the forwarding plane of the receiving IP Flex-Algorithm prefix capable routers when they participate in the associated topology and algorithm. Forwarding entries for IP Flex-Algorithm prefixes associated with Flex-Algorithms in which the node is not participating MUST NOT be installed in the forwarding plane.

When the IP Flex-Algorithm prefix is associated with a Flex-Algorithm, LFA paths to the prefix MUST be calculated using such Flex-Algorithm in the associated topology, to guarantee that they follow the same constraints as the calculation of the primary paths.

9. Deployment Considerations

IGP Flex-Algorithm can be used by many data-planes. The original specification was done for SR and SRv6, this specification adds IP as another data-plane that can use IGP Flex-Algorithm. Other data-planes may be defined in the future. This section provides some

details about the coexistence of the various data-planes of an IGP Flex-Algorithm.

Flex-Algorithm definition (FAD), as described in [I-D.ietf-lsr-flex-algo], is data-plane independent and is used by all Flex-Algorithm data-planes.

Participation in the Flex-Algorithm, as described in [I-D.ietf-lsr-flex-algo], is data-plane specific.

Calculation of the flex-algo paths is data-plane specific and uses data-plane specific participation advertisements.

Data-plane specific participation and calculation guarantee that the forwarding of the traffic over the Flex-Algorithm data-plane specific paths is consistent between all nodes that apply the IGP Flex-Algorithm to the data-plane.

Multiple data-planes can use the same Flex-Algorithm value at the same time and, and as such, share the FAD for it. For example, SR-MPLS and IP can both use a common Flex-Algorithm. Traffic for SR-MPLS will be forwarded based on Flex-algorithm specific SR SIDs. Traffic for IP Flex-Algorithm will be forwarded based on Flex-Algorithm specific prefix reachability advertisements. Note that for a particular Flex-Algorithm, for a particular IP prefix, there will only be path(s) calculated and installed for a single data-plane.

10. Protection

In many networks where IGP Flexible Algorithms are deployed, IGP restoration will be fast and additional protection mechanisms will not be required. IGP restoration may be enhanced by Equal Cost Multipath (ECMP).

In other networks, operators can deploy additional protection mechanisms. The following are examples:

- Loop Free Alternates (LFA) [RFC5286]
- Remote Loop Free Alternates (R-LFA) [RFC7490]

LFA and R-LFA computations MUST be restricted to the flex-algo topology and the computed backup nexthops should be programmed for the IP flex-algo prefixes.

11. IANA Considerations

This specification updates the OSPF Router Information (RI) TLVs Registry as follows:

Value	TLV Name	Reference
TBD	IP Algorithm TLV	This Document Section 5.2

Table 1

This document also updates the ISIS "Sub-TLVs for TLV 242" registry as follows:

Value	TLV Name	Reference
29	IP Algorithm Sub-TLV	This Document Section 5.1

Table 2

This document also updates the "IS-IS TLV Codepoints Registry" registry as follows:

Value	TLV Name	IIH	LSP	SNP	Purge	Reference
126	IPv4 Algorithm Prefix Reachability TLV	N	Y	N	N	This document, Section 6.1
127	IPv6 Algorithm Prefix Reachability TLV	N	Y	N	N	This document, Section 6.2

Table 3

The above TLVs share the sub-TLV space defined in "IS-IS Sub-TLVs for TLVs Advertising Prefix Reachability". This document updates the description of that registry by including IPv4 Algorithm Prefix Reachability TLV and IPv6 Algorithm Prefix Reachability TLV. It also includes these TLVs in a table which lists the presence of Sub-TLVs in a parent TLVs as follows:

Type	Description	126	127
1	32-bit Administrative Tag Sub-TLV	y	y
2	64-bit Administrative Tag Sub-TLV	y	y
3	Prefix Segment Identifier	n	n
4	Prefix Attribute Flags	y	y
5	SRv6 End SID	n	n
6	Flex-Algorithm Prefix Metric	n	n
11	IPv4 Source Router ID	y	y
12	IPv6 Source Router ID	y	y
32	BIER Info	n	n

This document updates the "OSPFv2 Extended Prefix TLV Sub-TLVs" registry as follows:

Value	TLV Name	Reference
TBD	OSPFv2 IP Algorithm Prefix Reachability TLV	This Document, Section 6.3

Table 4

This document updates the "OSPFv3 Extended-LSA Sub-TLVs" registry as follows:

Value	TLV Name	Reference
TBD	OSPFv3 IP Algorithm Prefix Reachability Sub-TLV	This Document, Section 6.4
TBD	OSPFv3 IP Flexible Algorithm ASBR Metric Sub-TLV	This Document, Section 6.5

Table 5

This document updates the "OSPFv2 Extended Inter-Area ASBR Sub-TLVs" registry as follows:

Value	TLV Name	Reference
2	OSPF IP Flexible Algorithm ASBR Metric Sub-TLV	This Document, Section 6.5

Table 6

12. Security Considerations

This document inherits security considerations from [I-D.ietf-lsr-flex-algo].

13. Acknowledgements

Thanks to Bruno Decraene for his contributions to this document. Special thanks to Petr Bonbon Adamec of Cesnet for supporting interoperability testing.

14. References

14.1. Normative References

[I-D.ietf-lsr-flex-algo]
Psenak, P., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", draft-ietf-lsr-flex-algo-19 (work in progress), April 2022.

[ISO10589]
IANA, "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>.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <https://www.rfc-editor.org/info/rfc791>.

[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>.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, DOI 10.17487/RFC3630, September 2003, <https://www.rfc-editor.org/info/rfc3630>.

[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>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <https://www.rfc-editor.org/info/rfc5305>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", RFC 7684, DOI 10.17487/RFC7684, November 2015, <https://www.rfc-editor.org/info/rfc7684>.

[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>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[RFC8362] Lindem, A., Roy, A., Goethals, D., Reddy Vallem, V., and F. Baker, "OSPFv3 Link State Advertisement (LSA) Extensibility", RFC 8362, DOI 10.17487/RFC8362, April 2018, <https://www.rfc-editor.org/info/rfc8362>.

14.2. Informative References

[IANA-ALG] IANA, "Sub-TLVs for TLV 242 (IS-IS Router CAPABILITY TLV)", August 1987, <https://www.iana.org/assignments/igp-parameters/igp-parameters.xhtml#igp-algorithm-types>.

[RFC5286] Atlas, A., Ed. and A. Zinin, Ed., "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, DOI 10.17487/RFC5286, September 2008, <https://www.rfc-editor.org/info/rfc5286>.

[RFC7490] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, DOI 10.17487/RFC7490, April 2015, <https://www.rfc-editor.org/info/rfc7490>.

[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>.

Authors' Addresses

William Britto
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore, Karnataka 560103
India

Email: bwilliam@juniper.net

Shraddha Hegde
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore, Karnataka 560103
India

Email: shraddha@juniper.net

Parag Kaneriya
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore, Karnataka 560103
India

Email: pkaneria@juniper.net

Rejesh Shetty
Juniper Networks
Elnath-Exora Business Park Survey
Bangalore, Karnataka 560103
India

Email: mrajesh@juniper.net

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, Virginia 20171
USA

Email: rbonica@juniper.net

Peter Psenak
Cisco Systems
Apollo Business Center
Mlynske nivy 43, Bratislava 82109
Slovakia

Email: ppsenak@cisco.com