

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 14, 2019

P. Psenak, Ed.
Cisco Systems
S. Hegde
Juniper Networks, Inc.
C. Filsfils
K. Talaulikar
Cisco Systems, Inc.
A. Gulko
Thomson Reuters
May 13, 2019

IGP Flexible Algorithm
draft-ietf-lsr-flex-algo-02.txt

Abstract

IGP protocols traditionally compute best paths over the network based on the IGP metric assigned to the links. Many network deployments use RSVP-TE based or Segment Routing based Traffic Engineering to enforce traffic over a path that is computed using different metrics or constraints than the shortest IGP path. This document proposes a solution that allows IGPs themselves to compute constraint based paths over the network. This document also specifies a way of using Segment Routing Prefix-SIDs to steer packets along the constraint-based paths.

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 14, 2019.

Copyright Notice

Copyright (c) 2019 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Requirements notation	4
3.	Terminology	4
4.	Flexible Algorithm	4
5.	Flexible Algorithm Definition Advertisement	5
5.1.	ISIS Flexible Algorithm Definition Sub-TLV	5
5.2.	OSPF Flexible Algorithm Definition TLV	7
5.3.	Common Handling of Flexible Algorithm Definition TLV	8
6.	Sub-TLVs of ISIS FAD Sub-TLV	9
6.1.	ISIS Flexible Algorithm Exclude Admin Group Sub-TLV	9
6.2.	ISIS Flexible Algorithm Include-Any Admin Group Sub-TLV	10
6.3.	ISIS Flexible Algorithm Include-All Admin Group Sub-TLV	10
7.	Sub-TLVs of OSPF FAD TLV	11
7.1.	OSPF Flexible Algorithm Exclude Admin Group Sub-TLV	11
7.2.	OSPF Flexible Algorithm Include-Any Admin Group Sub-TLV	11
7.3.	OSPF Flexible Algorithm Include-All Admin Group Sub-TLV	12
8.	Advertisement of Node Participation in a Flex-Algorithm	12
8.1.	Advertisement of Node Participation for Segment Routing	12
8.2.	Advertisement of Node Participation for Other Applications	13
9.	Advertisement of Link Attributes for Flex-Algorithm	13
10.	Calculation of Flexible Algorithm Paths	13
11.	Flex-Algorithm and Forwarding Plane	15
11.1.	Segment Routing MPLS Forwarding for Flex-Algorithm	15
11.2.	Other Applications' Forwarding for Flex-Algorithm	16
12.	Backward Compatibility	16
13.	Security Considerations	16
14.	IANA Considerations	16
14.1.	IGP IANA Considerations	16
14.1.1.	IGP Algorithm Types Registry	16
14.1.2.	Flexible Algorithm Definition Metric-Type Registry	17

14.2.	ISIS IANA Considerations	17
14.2.1.	Sub TLVs for Type 242	17
14.2.2.	Sub-Sub-TLVs for Flexible Algorithm Definition Sub-TLV	17
14.3.	OSPF IANA Considerations	18
14.3.1.	OSPF Router Information (RI) TLVs Registry	18
14.3.2.	OSPF Flexible Algorithm Definition TLV Sub-TLV Registry	18
15.	Contributors	19
16.	References	20
16.1.	Normative References	20
16.2.	Informative References	21
	Authors' Addresses	22

[1. Introduction](#)

An IGP computed path based on the shortest IGP metric must often be replaced by traffic engineered path due to the traffic requirements which are not reflected by the IGP metric. Some networks engineer the IGP metric assignments in a way that the IGP Metric reflects the link bandwidth or delay. If, for example, the IGP metric is reflecting the bandwidth on the link and the application traffic is delay sensitive, the best IGP path may not reflect the best path from such application's perspective.

To overcome this limitation, various sorts of traffic engineering have been deployed, including RSVP-TE and SR-TE, in which case the TE component is responsible for computing the path based on additional metrics and/or constraints. Such paths need to be installed in the forwarding tables in addition to, or as a replacement for the original paths computed by IGPs. Tunnels are often used to represent the engineered paths and mechanisms like one described in [[RFC3906](#)] are used to replace the native IGP paths with such tunnel paths.

This document specifies a set of extensions to ISIS, OSPFv2 and OSPFv3 that enable a router to send TLVs that (a) describe a set of constraints on the topology, (b) identify calculation-type, and (c) metric-type that are to be used to compute the best paths along the constrained topology. A given combination of calculation-type, metric-type and constraints is known as a "Flexible Algorithm Definition". A router that sends such a set of TLVs also assigns a specific value, Flex-Algorithm, to the specified combination of calculation-type, metric-type and constraints.

This document also specifies a way for a router to use IGPs to associate one or more Segment Routing Prefix-SIDs with a particular Flex-Algorithm. Each such Prefix-SID then represents a path that is computed according to the identified Flex-Algorithm.

2. Requirements notation

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 [\[BCP14\]](#) [\[RFC2119\]](#) [\[RFC8174\]](#) when, and only when, they appear in all capitals, as shown here.

3. Terminology

This section defines terms that are often used in this document.

Flexible Algorithm Definition - the set consisting of (a) calculation-type, (b) metric-type and (c) a set of constraints.

Flexible Algorithm - a numeric identifier in the range 128-255 that is associated via provisioning with the Flexible-Algorithm Definition.

Local Flexible Algorithm Definition - Flexible Algorithm Definition defined locally on the node.

Remote Flexible Algorithm Definition - Flexible Algorithm Definition received from other nodes via IGP flooding.

Flexible Algorithm Participation - per application configuration state that expresses whether the node is participating in a particular Flexible Algorithm.

IGP Algorithm - value from the the "IGP Algorithm Types" registry defined under "Interior Gateway Protocol (IGP) Parameters" IANA registries. IGP Algorithms represents the triplet (Calculation Type, Metric, Constraints), where the second and third elements of the triple MAY not exist.

4. Flexible Algorithm

Many possible constraints may be used to compute a path over a network. Some networks are deployed as multiple planes. A simple form of constraint may be to use a particular plane. A more sophisticated form of constraint can include some extended metric as described in [\[RFC7810\]](#). Constraints which restrict paths to links with specific affinities or avoid links with specific affinities are also possible. Combinations of these are also possible.

To provide maximum flexibility, we want to provide a mechanism that allows a router to (a) identify a particular calculation-type, (b) metric-type, (c) describe a particular set of constraints, and (d)

assign a numeric identifier, referred to as Flex-Algorithm, to the combination of that calculation-type, metric-type and those constraints. We want the mapping between the Flex-Algorithm and it's meaning to be flexible and defined by the user. As long as all routers in the domain have a common understanding as to what a particular Flex-Algorithm represents, the resulting routing computation is consistent and traffic is not subject to any looping.

The set consisting of (a) calculation-type, (b) metric-type and (c) a set of constraints is referred to as a Flexible-Algorithm Definition.

Flexible-Algorithm is a numeric identifier in the range 128-255 that is associated via provisioning with the Flexible-Algorithm Definition.

IANA "IGP Algorithm Types" registry defines the set of values for IGP Algorithms. We propose to allocate the following values for Flex-Algorithms from this registry:

128-255 - Flex-Algorithms

5. Flexible Algorithm Definition Advertisement

To guarantee the loop free forwarding for paths computed for a particular Flex-Algorithm, all routers that (a) are configured to participate in a particular Flex-Algorithm, and (b) are in the same Flex-Algorithm definition advertisement scope MUST agree on the definition of the Flex-Algorithm.

5.1. ISIS Flexible Algorithm Definition Sub-TLV

ISIS Flexible Algorithm Definition Sub-TLV (FAD Sub-TLV) is used to advertise the definition of the Flex-Algorithm.

ISIS FAD Sub-TLV is advertised as a Sub-TLV of the ISIS Router Capability TLV-242 that is defined in [[RFC7981](#)].

ISIS FAD Sub-TLV 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      |Flex-Algorithm |  Metric-Type  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Calc-Type    |      Priority    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Sub-TLVs                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     ...                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Type: TBD, suggested value 26

Length: variable, dependent on the included Sub-TLVs

Flex-Algorithm: Single octet value between 128 and 255 inclusive.

Metric-Type: Type of metric to be used during the calculation.
Following values are defined:

0: IGP Metric

1: Min Unidirectional Link Delay as defined in [[RFC7810](#)].

2: TE default metric as defined in [[RFC5305](#)].

Calc-Type: value from 0 to 127 inclusive from the "IGP Algorithm Types" registry defined under "Interior Gateway Protocol (IGP) Parameters" IANA registries. IGP algorithms in the range of 0-127 have a defined triplet (Calculation Type, Metric, Constraints). When used to specify the Calc-Type in the FAD Sub-TLV, only the Calculation Type defined for the specified IGP Algorithm is used. The Metric/Constraints MUST NOT be inherited. If the required calculation type is Shortest Path First, the value 0 SHOULD appear in this field.

Priority: Value between 0 and 255 inclusive that specifies the priority of the advertisement.

Sub-TLVs - optional sub-TLVs.

The ISIS FAD Sub-TLV MAY be flooded only in a given level or throughout the domain. In the latter case the S-flag is set as

When multiple OSPF FAD TLVs, for the same Flexible-Algorithm, are received from a given router, the receiver MUST use the first occurrence of the TLV in the Router Information LSA. If the OSPF FAD TLV, for the same Flex-Algorithm, appears in multiple Router Information LSAs that have different flooding scopes, the OSPF FAD TLV in the Router Information LSA with the area-scoped flooding scope

MUST be used. If the OSPF FAD TLV, for the same algorithm, appears in multiple Router Information LSAs that have the same flooding scope, the OSPF FAD TLV in the Router Information (RI) LSA with the numerically smallest Instance ID MUST be used and subsequent instances of the OSPF FAD TLV MUST be ignored.

The RI LSA can be advertised at any of the defined opaque flooding scopes (link, area, or Autonomous System (AS)). For the purpose of OSPF FAD TLV advertisement, area-scoped flooding is REQUIRED. The Autonomous System flooding scope SHOULD not be used by default unless local configuration policy on the originating router indicates domain wide flooding.

5.3. Common Handling of Flexible Algorithm Definition TLV

This section describes the protocol independent handling of the FAD TLV (OSPF) or FAD Sub-TLV (ISIS). We will refer to it as FAD TLV in this section, even though in case of ISIS it is a Sub-TLV.

The value of the Flex-Algorithm MUST be between 128 and 255 inclusive. If it is not, the FAD TLV MUST be ignored.

Not every router configured to participate in a particular Flex-Algorithm need a local definition of such Flex-Algorithm. Only a subset of the routers participating in the particular Flex-Algorithm need the local definition of the Flex-Algorithm.

Every router, that is configured to participate in a particular Flex-Algorithm, MUST select the Flex-Algorithm definition based on the following ordered rules. This allows for the consistent Flex-Algorithm definition selection in cases where different routers advertise different definitions for a given Flex-Algorithm:

1. From the advertisements of the FAD in the area (including both locally generated advertisements and received advertisements) select the one(s) with the highest priority.
2. If there are multiple advertisements of the FAD with the same highest priority, select the one that is originated from the router with the highest System-ID in case of ISIS or Router ID in case of OSPFv2 and OSPFv3. For ISIS the System-ID is described in [ISO10589]. For OSPFv2 and OSPFv3 standard Router ID is described in [RFC2328] and [RFC5340] respectively.

A router that is not configured to participate in a particular Flex-Algorithm MUST ignore FAD Sub-TLVs advertisements for such Flex-Algorithm.

Any change in the Flex-Algorithm definition may result in temporary disruption of traffic that is forwarded based on such Flex-Algorithm paths. The impact is similar to any other event that requires network wide convergence.

If a node is configured to participate in a particular Flexible-Algorithm, but the selected Flex-Algorithm definition includes calculation-type, metric-type or constraint that is not supported by the node, it **MUST** stop participating in such Flexible-Algorithm. That implies that it **MUST NOT** announce participation for such Flexible-Algorithm and it **MUST** remove any forwarding state associated with it.

Flex-Algorithm definition is topology independent. It applies to all topologies that a router participates in.

6. Sub-TLVs of ISIS FAD Sub-TLV

6.1. ISIS Flexible Algorithm Exclude Admin Group Sub-TLV

The Flexible-Algorithm definition can specify 'colors' that are used by the operator to exclude links during the Flex-Algorithm path computation.

Flexible Algorithm Exclude Admin Group Sub-TLV (FAEAG Sub-TLV) is a Sub-TLV of the ISIS FAD Sub-TLV. It 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      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Extended Admin Group                                     |
+-+-----+
|                                     ...                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Type: 1

Length: variable, dependent on the size of the Extended Admin Group. **MUST** be a multiple of 4 octets.

Extended Administrative Group: Extended Administrative Group as defined in [[RFC7308](#)].

ISIS FAEAG Sub-TLV MAY NOT appear more than once in an ISIS FAD Sub-TLV. If it appears more than once, the ISIS FAD Sub-TLV MUST be ignored by the receiver.

6.2. ISIS Flexible Algorithm Include-Any Admin Group Sub-TLV

The Flexible-Algorithm definition can specify 'colors' that are used by the operator to include link during the Flex-Algorithm path computation.

ISIS Flexible Algorithm Include-Any Admin Group Sub-TLV is used to advertise include-any rule that is used during the Flex-Algorithm path calculation as specified in Section [Section 10](#).

The format of the ISIS Flexible Algorithm Include-Any Admin Group Sub-TLV is identical to the format of the FAEAG Sub-TLV in [Section 6.1](#).

Flexible Algorithm Include-Any Admin Group Sub-TLV Type is 2.

ISIS Flexible Algorithm Include-Any Admin Group Sub-TLV MAY NOT appear more than once in an ISIS FAD Sub-TLV. If it appears more than once, the ISIS FAD Sub-TLV MUST be ignored by the receiver.

6.3. ISIS Flexible Algorithm Include-All Admin Group Sub-TLV

The Flexible-Algorithm definition can specify 'colors' that are used by the operator to include link during the Flex-Algorithm path computation.

ISIS Flexible Algorithm Include-All Admin Group Sub-TLV is used to advertise include-all rule that is used during the Flex-Algorithm path calculation as specified in Section [Section 10](#).

The format of the ISIS Flexible Algorithm Include-All Admin Group Sub-TLV is identical to the format of the FAEAG Sub-TLV in [Section 6.1](#).

ISIS Flexible Algorithm Include-All Admin Group Sub-TLV Type is 3.

ISIS Flexible Algorithm Include-All Admin Group Sub-TLV MAY NOT appear more than once in an ISIS FAD Sub-TLV. If it appears more than once, the ISIS FAD Sub-TLV MUST be ignored by the receiver.

7. Sub-TLVs of OSPF FAD TLV

7.1. OSPF Flexible Algorithm Exclude Admin Group Sub-TLV

Flexible Algorithm Exclude Admin Group Sub-TLV (FAEAG Sub-TLV) is a Sub-TLV of the OSPF FAD TLV. It's usage is described in [Section 6.1](#). It 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                 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
|      Extended Admin Group     |                               |
+-+-----+-----+-----+-----+-----+-----+-----+---+
|                               |                               |
|                               |      ...                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Type: 1

Length: variable, dependent on the size of the Extended Admin Group. MUST be a multiple of 4 octets.

Extended Administrative Group: Extended Administrative Group as defined in [[RFC7308](#)].

OSPF FAEAG Sub-TLV MAY NOT appear more than once in an OSPF FAD TLV. If it appears more than once, the OSPF FAD TLV MUST be ignored by the receiver.

7.2. OSPF Flexible Algorithm Include-Any Admin Group Sub-TLV

The usage of this Sub-TLVs is described in [Section 6.2](#).

The format of the OSPF Flexible Algorithm Include-Any Admin Group Sub-TLV is identical to the format of the OSPF FAEAG Sub-TLV in [Section 7.1](#).

Flexible Algorithm Include-Any Admin Group Sub-TLV Type is 2.

OSPF Flexible Algorithm Include-Any Admin Group Sub-TLV MAY NOT appear more than once in an OPSF FAD TLV. If it appears more than once, the OSPF FAD TLV MUST be ignored by the receiver.

7.3. OSPF Flexible Algorithm Include-All Admin Group Sub-TLV

The usage of this Sub-TLVs is described in [Section 6.3](#).

The format of the OSPF Flexible Algorithm Include-Any Admin Group Sub-TLV is identical to the format of the OSPF FAEAG Sub-TLV in [Section 7.1](#).

Flexible Algorithm Include-Any Admin Group Sub-TLV Type is 3.

OSPF Flexible Algorithm Include-All Admin Group Sub-TLV MAY NOT appear more than once in an OSPF FAD TLV. If it appears more than once, the OSPF FAD TLV MUST be ignored by the receiver.

8. Advertisement of Node Participation in a Flex-Algorithm

When a router is configured to support a particular Flex-Algorithm, we say it is participating in that Flex-Algorithm.

Paths computed for a specific Flex-Algorithm MAY be used by various applications, each potentially using its own specific data plane for forwarding the data over such paths. To guarantee the presence of the application specific forwarding state associated with a particular Flex-Algorithm, a router MUST advertise its participation for a particular Flex-Algorithm for each application specifically.

8.1. Advertisement of Node Participation for Segment Routing

[[I-D.ietf-isis-segment-routing-extensions](#)], [[I-D.ietf-ospf-segment-routing-extensions](#)] and [[I-D.ietf-ospf-ospfv3-segment-routing-extensions](#)] (IGP Segment Routing extensions) describe how SR-Algorithm is used to define how the best path is computed by the IGP.

Routers advertise the support for the SR-Algorithm as a node capability as described in the above mentioned IGP Segment Routing extensions. To advertise participation for a particular Flex-Algorithm for Segment Routing, the Flex-Algorithm value MUST be advertised in the SR-Algorithm TLV (OSPF) or sub-TLV (ISIS).

Segment Routing Flex-Algorithm participation advertisement is topology independent. When a router advertises participation in an SR-Algorithm, the participation applies to all topologies in which the advertising node participates.

8.2. Advertisement of Node Participation for Other Applications

This section describes considerations related to how other applications can advertise its participation in a specific Flex-Algorithm.

Application specific Flex-Algorithm participation advertisements MAY be topology specific or MAY be topology independent, depending on the application itself.

Application specific advertisement for Flex-Algorithm participation MUST be defined for each application and is outside of the scope of this document.

9. Advertisement of Link Attributes for Flex-Algorithm

Various link include or exclude rules can be part of the Flex-Algorithm definition. These rules use Admin Groups (AG) as defined in [[RFC7308](#)] and [[RFC5305](#)], or Extended Administrative Groups (EAG) as defined in [[RFC7308](#)].

To advertise a link affinity in a form of the AG or EAG that is used during Flex-Algorithm calculation, an Application Specific Link Attributes sub-TLV as described in [[I-D.ietf-isis-te-app](#)], or sub-TLV of Extended Link TLV as described in [[I-D.ietf-ospf-te-link-attr-reuse](#)] MUST be used. The advertisement MUST indicate that it is usable by the Flex-Algorithm application.

10. Calculation of Flexible Algorithm Paths

A router MUST be configured to participate in a given Flex-Algorithm K before it can compute any path for that Flex-Algorithm.

A router which participates in a given Flex Algorithm MUST use the FAD selected based on the rules defined in Section [Section 5.3](#).

As described in [Section 8](#), participation for any particular Flex-Algorithm MUST be advertised on a per application basis. Calculation of the paths for any particular Flex-Algorithm MUST be application specific.

The way applications handle nodes that do not participate in Flexible-Algorithm is application specific. If the application only wants to consider participating nodes during the Flex-Algorithm calculation, then when computing paths for a given Flex-Algorithm, all nodes that do not advertise participation for that Flex-Algorithm in the application specific advertisements MUST be pruned from the

topology. MPLS Segment Routing is an application that MUST use such pruning when computing Flex-Algorithm paths.

When computing the path for a give Flex-Algorithm, the metric-type that is part of the Flex-Algorithm definition ([Section 5](#)) MUST be used.

When computing the path for a given Flex-Algorithm, the calculation-type that is part of the Flex-Algorithm definition ([Section 5](#)) MUST be used.

Various link include or exclude rules can be part of the Flex-Algorithm definition. To refer to particular bit within an AG or EAG we uses term 'color'.

Rules, in the order as specified below, MUST be used to prune link from the topology during the Flex-Algorithm computation.

For all links in the topology:

1. Check if any exclude rule is part of the Flex-Algorithm definition. If such exclude rule exists, check if any color that is part of the exclude rule is also set on the link. If such a color exist, the link MUST be pruned from the computation.
2. Check if any include-any rule is part of the Flex-Algorithm definition. if such include-any rule exists, check if any color that is part of the include-any rule is also set on the link. If such color does not exist, the link MUST be pruned from the computation.
3. Check if any include-all rule is part of the Flex-Algorithm definition. If such include-all rule exists, check if all colors that are part of the include-all rule are also set on the link. If not all such colors are set on the link, the link MUST be pruned from the computation.
4. If the Flex-Algorithm definition uses other than IGP metric ([Section 5](#)), and such metric is not advertised for the particular link in a topology for which the computation is done, such link MUST be pruned from the computation. A metric of value 0 MUST NOT be assumed in such case.

Any IGP Shortest Path Tree calculation is limited to a single area. Same applies to Flex-Algorithm calculations. Given that the computing router may not have the visibility to the topology of remote areas, the Flex-Algorithm specific path to an inter-area prefix will only be computed for the local area only. The egress L1/

L2 router (ABR in OSPF) will be selected based on the best path for the given Flex-Algorithm in the local area and such egress L1/L2 (ABR in OSPF) router will be responsible to compute the best Flex-Algorithm specific path over the next area. This may produce an end-to-end path, which is sub-optimal based on Flex-Algorithm constraints. If the best end-to-end path for a given Flex-Algorithm needs to be used for inter-area destinations, paths for such destinations need to be computed by the entity that has the topological information about all areas.

11. Flex-Algorithm and Forwarding Plane

This section describes how Flex-Algorithm paths are used with forwarding.

11.1. Segment Routing MPLS Forwarding for Flex-Algorithm

This section describes how Flex-Algorithm paths are used with SR MPLS forwarding.

Prefix SID advertisements include an SR-Algorithm value and as such are associated with the specified SR-Algorithm. Prefix-SIDs are also associated with a specific topology which is inherited from the associated prefix reachability advertisement. When the algorithm value advertised is a Flex-Algorithm value, the Prefix SID is associated with paths calculated using that Flex-Algorithm in the associated topology.

A Flex-Algorithm path **MUST** be installed in the MPLS forwarding plane using the MPLS label that corresponds to the Prefix-SID that was advertised for that Flex-algorithm. If the Prefix SID for a given Flex-algorithm is not known, the Flex-Algorithm specific path cannot be installed in the MPLS forwarding plane.

Traffic that is supposed to be routed via Flex-Algorithm specific paths, **MUST** be dropped where there are no such paths available.

Loop Free Alternate (LFA) paths for a given Flex-Algorithm **MUST** be computed using the same constraints as the calculation of the primary paths for that Flex-Algorithm. LFA paths **MUST** only use Prefix-SIDs advertised specifically for the given algorithm. LFA paths **MUST NOT** use an Adjacency-SID that belongs to a link that has been pruned from the Flex-Algorithm computation.

If LFA protection is being used to protect a given Flex-Algorithm paths, all routers in the area participating in the given Flex-Algorithm **SHOULD** advertise at least one Flex-Algorithm specific Node-

SID. These Node-SIDs are used to enforce traffic over the LFA computed backup path.

11.2. Other Applications' Forwarding for Flex-Algorithm

Any application that wants to use Flex-Algorithm specific forwarding need to install some form of Flex-Algorithm specific forwarding entries.

Application specific forwarding for Flex-Algorithm MUST be defined for each application and is outside of the scope of this document.

12. Backward Compatibility

This extension brings no new backward compatibility issues.

13. Security Considerations

This draft adds a two new ways to disrupt the IGP networks:

An attacker can hijack a particular Flex-Algorithm by advertising a FAD with a priority of 255 (or any priority higher than that of the legitimate nodes).

An attacker could make it look like a router supports a particular Flex-Algorithm when it actually doesn't, or vice versa.

Both of these attacks can be addressed by the existing security extensions as described in [[RFC5304](#)] and [[RFC5310](#)] for ISIS, in [[RFC2328](#)] and [[RFC7474](#)] for OSPFv2 and in [[RFC5340](#)] and [[RFC4552](#)] for OSPFv3.

14. IANA Considerations

14.1. IGP IANA Considerations

14.1.1. IGP Algorithm Types Registry

This document makes the following registrations in the "IGP Algorithm Types" registry:

Type: 128-255.

Description: Flexible Algorithms.

Reference: This document ([Section 4](#)).

14.1.2. Flexible Algorithm Definition Metric-Type Registry

IANA is requested to set up a registry called "Flexible Algorithm Definition Metric-Type Registry" under a "Interior Gateway Protocol (IGP) Parameters" IANA registries. The registration policy for this registry is "Standards Action" ([[RFC8126](#)] and [[RFC7120](#)]).

Values in this registry come from the range 0-255.

This document registers following values in the "Flexible Algorithm Definition Metric-Type Registry":

Type: 0

Description: IGP metric

Reference: This document ([Section 5.1](#))

Type: 1

Description: Min Unidirectional Link Delay [[RFC7810](#)]

Reference: This document ([Section 5.1](#))

Type: 2

Description: TE Default Metric [[RFC5305](#)]

Reference: This document ([Section 5.1](#))

14.2. ISIS IANA Considerations

14.2.1. Sub TLVs for Type 242

This document makes the following registrations in the "sub-TLVs for TLV 242" registry.

Type: TBD (suggested value 26).

Description: Flexible Algorithm Definition Sub-TLV.

Reference: This document ([Section 5.1](#)).

14.2.2. Sub-Sub-TLVs for Flexible Algorithm Definition Sub-TLV

This document creates the following Sub-Sub-TLV Registry:

Registry: Sub-Sub-TLVs for Flexible Algorithm Definition Sub-TLV

Registration Procedure: Expert review

Reference: This document ([Section 5.1](#))

This document defines the following Sub-Sub-TLVs in the "Sub-Sub-TLVs for Flexible Algorithm Definition Sub-TLV" registry:

Type: 1

Description: Flexible Algorithm Exclude Admin Group Sub-TLV

Reference: This document ([Section 6.1](#)).

Type: 2

Description: Flexible Algorithm Include-Any Admin Group Sub-TLV

Reference: This document ([Section 6.2](#)).

Type: 3

Description: Flexible Algorithm Include-All Admin Group Sub-TLV

Reference: This document ([Section 6.3](#)).

[14.3.](#) OSPF IANA Considerations

[14.3.1.](#) OSPF Router Information (RI) TLVs Registry

This specification updates the OSPF Router Information (RI) TLVs Registry with the following value:

- o TBD (suggested value 16) - Flexible Algorithm Definition TLV

[14.3.2.](#) OSPF Flexible Algorithm Definition TLV Sub-TLV Registry

This document creates the following registry:

Registry: OSPF Flexible Algorithm Definition TLV sub-TLV

Registration Procedure: Expert review

Reference: This document ([Section 5.2](#))

The "OSPF Flexible Algorithm Definition TLV sub-TLV" registry will define sub-TLVs at any level of nesting for Flexible Algorithm TLV and should be added to the "Open Shortest Path First (OSPF)

Parameters" registries group. New values can be allocated via IETF Review or IESG Approval.

This document registers following Sub-TLVs in the "TLVs for Flexible Algorithm Definition TLV" registry:

Type: 1

Description: Flexible Algorithm Exclude Admin Group Sub-TLV

Reference: This document ([Section 7.1](#)).

Type: 2

Description: Flexible Algorithm Include-Any Admin Group Sub-TLV

Reference: This document ([Section 7.2](#)).

Type: 3

Description: Flexible Algorithm Include-All Admin Group Sub-TLV

Reference: This document ([Section 7.3](#)).

Types in the range 32768-33023 are for experimental use; these will not be registered with IANA, and MUST NOT be mentioned by RFCs.

Types in the range 33024-65535 are not to be assigned at this time. Before any assignments can be made in the 33024-65535 range, there MUST be an IETF specification that specifies IANA Considerations that covers the range being assigned.

15. Contributors

This draft, among other things, is also addressing the problem that the [[I-D.gulkohegde-routing-planes-using-sr](#)] was trying to solve. All authors of that draft agreed to join this draft.

Thanks to Eric Rosen, Les Ginsberg and Tony Przygienda for their detailed review and excellent comments.

Thanks to Cengiz Halit for his review and feedback during initial phase of the solution definition.

Thanks to Kenji Kumaki for his comments.

16. References

16.1. Normative References

- [BCP14] , <<https://tools.ietf.org/html/bcp14>>.
- [I-D.ietf-isis-segment-routing-extensions]
Previdi, S., Ginsberg, L., Filsfils, C., Bashandy, A., Gredler, H., and B. Decraene, "IS-IS Extensions for Segment Routing", [draft-ietf-isis-segment-routing-extensions-24](#) (work in progress), April 2019.
- [I-D.ietf-isis-te-app]
Ginsberg, L., Psenak, P., Previdi, S., Henderickx, W., and J. Drake, "IS-IS TE Attributes per application", [draft-ietf-isis-te-app-06](#) (work in progress), April 2019.
- [I-D.ietf-ospf-ospfv3-segment-routing-extensions]
Psenak, P. and S. Previdi, "OSPFv3 Extensions for Segment Routing", [draft-ietf-ospf-ospfv3-segment-routing-extensions-23](#) (work in progress), January 2019.
- [I-D.ietf-ospf-segment-routing-extensions]
Psenak, P., Previdi, S., Filsfils, C., Gredler, H., Shakir, R., Henderickx, W., and J. Tantsura, "OSPF Extensions for Segment Routing", [draft-ietf-ospf-segment-routing-extensions-27](#) (work in progress), December 2018.
- [I-D.ietf-ospf-te-link-attr-reuse]
Psenak, P., Ginsberg, L., Henderickx, W., Tantsura, J., and J. Drake, "OSPF Link Traffic Engineering (TE) Attribute Reuse", [draft-ietf-ospf-te-link-attr-reuse-07](#) (work in progress), April 2019.
- [ISO10589]
International Organization for Standardization, "Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, Nov 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>>.

- [RFC7308] Osborne, E., "Extended Administrative Groups in MPLS Traffic Engineering (MPLS-TE)", [RFC 7308](#), DOI 10.17487/RFC7308, July 2014, <<https://www.rfc-editor.org/info/rfc7308>>.
- [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>>.

[16.2](#). Informative References

- [I-D.gulkohegde-routing-planes-using-sr] Hegde, S. and a. arkadiy.gulko@thomsonreuters.com, "Separating Routing Planes using Segment Routing", [draft-gulkohegde-routing-planes-using-sr-00](#) (work in progress), March 2017.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.
- [RFC3906] Shen, N. and H. Smit, "Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels", [RFC 3906](#), DOI 10.17487/RFC3906, October 2004, <<https://www.rfc-editor.org/info/rfc3906>>.
- [RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", [RFC 4552](#), DOI 10.17487/RFC4552, June 2006, <<https://www.rfc-editor.org/info/rfc4552>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", [RFC 5304](#), DOI 10.17487/RFC5304, October 2008, <<https://www.rfc-editor.org/info/rfc5304>>.
- [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>>.

- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", [RFC 5310](#), DOI 10.17487/RFC5310, February 2009, <<https://www.rfc-editor.org/info/rfc5310>>.
- [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>>.
- [RFC7120] Cotton, M., "Early IANA Allocation of Standards Track Code Points", [BCP 100](#), [RFC 7120](#), DOI 10.17487/RFC7120, January 2014, <<https://www.rfc-editor.org/info/rfc7120>>.
- [RFC7474] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed., "Security Extension for OSPFv2 When Using Manual Key Management", [RFC 7474](#), DOI 10.17487/RFC7474, April 2015, <<https://www.rfc-editor.org/info/rfc7474>>.
- [RFC7810] Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", [RFC 7810](#), DOI 10.17487/RFC7810, May 2016, <<https://www.rfc-editor.org/info/rfc7810>>.
- [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>>.

Authors' Addresses

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

Email: ppsenak@cisco.com

Shraddha Hegde
Juniper Networks, Inc.
Embassy Business Park
Bangalore, KA, 560093
India

Email: shraddha@juniper.net

Clarence Filsfils
Cisco Systems, Inc.
Brussels
Belgium

Email: cfilsfil@cisco.com

Ketan Talaulikar
Cisco Systems, Inc.
S.No. 154/6, Phase I, Hinjawadi
PUNE, MAHARASHTRA 411 057
India

Email: ketant@cisco.com

Arkadiy Gulko
Thomson Reuters

Email: arkadiy.gulko@thomsonreuters.com

