

SPRING
Internet-Draft
Intended status: Standards Track
Expires: September 9, 2021

S. Hegde
W. Britto
R. Shetty
Juniper Networks Inc.
B. Decraene
Orange
P. Psenak
Cisco Systems
T. Li
Arista Networks
March 8, 2021

Flexible Algorithms: Bandwidth, Delay, Metrics and Constraints
draft-hegde-lsr-flex-algo-bw-con-01

Abstract

Many networks configure the link metric relative to the link capacity. High bandwidth traffic gets routed as per the link capacity. Flexible algorithms provides mechanisms to create constraint based paths in IGP. This draft documents a set of bandwidth related constraints to be used in Flexible Algorithms.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

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 September 9, 2021.

Copyright Notice

Copyright (c) 2021 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 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.	Bandwidth Metric Advertisement	3
2.1.	ISIS Bandwidth Metric sub-TLV	4
2.2.	OSPF Bandwidth Metric sub-TLV	5
3.	FAD constraint sub-TLVs	5
3.1.	ISIS FAD constraint sub-TLVs	6
3.1.1.	ISIS Exclude Minimum Bandwidth sub-TLV	6
3.1.2.	ISIS Exclude Maximum Delay sub-TLV	7
3.2.	OSPF FAD constraint sub-TLVs	7
3.2.1.	OSPF Exclude Minimum Bandwidth sub-TLV	8
3.2.2.	OSPF Exclude Maximum Delay sub-TLV	8
4.	Automatic Metric Calculation	9
4.1.	Automatic Metric Calculation Modes	10
4.1.1.	Simple Mode	10
4.1.2.	Interface Group Mode	10
4.2.	Automatic Metric Calculation Methods	10
4.2.1.	Reference Bandwidth method	11
4.2.2.	Bandwidth Thresholds method	11
4.3.	ISIS FAD constraint sub-TLVs for automatic metric calculation	11
4.3.1.	Reference Bandwidth sub-TLV	11
4.3.2.	Bandwidth Thresholds sub-TLV	13
4.4.	OSPF FAD constraint sub-TLVs for automatic metric calculation	15
4.4.1.	Reference Bandwidth sub-TLV	15
4.4.2.	Bandwidth Threshold sub-TLV	17
5.	Bandwidth metric considerations	19
6.	Calculation of Flex-Algorithm paths	19
7.	Backward Compatibility	20
8.	Security Considerations	20
9.	IANA Considerations	20

9.1.	IGP Metric-Type Registry	20
9.2.	ISIS Sub-Sub-TLVs for Flexible Algorithm Definition Sub-TLV	20
9.3.	OSPF Sub-TLVs for Flexible Algorithm Definition Sub-TLV .	21
9.4.	Sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223	21
9.5.	Sub-sub-TLV Codepoints for Application-Specific Link Attributes	21
9.6.	OSPFv2 Extended Link TLV Sub-TLVs	21
9.7.	Types for sub-TLVs of TE Link TLV (Value 2)	22
10.	Acknowledgements	22
11.	Contributors	22
12.	References	22
12.1.	Normative References	22
12.2.	Informative References	23
	Authors' Addresses	23

[1.](#) Introduction

High bandwidth traffic such as residential internet traffic and machine to machine elephant flows benefit from using high capacity links for the traffic. Many network operators define link metric relative to the link capacity. It may be useful to exclude the high bandwidth traffic from utilizing links below certain capacity. A Flex-Algorithm [[I-D.ietf-lsr-flex-algo](#)] is defined as a set of parameters consisting of calculation-type, metric-type and a set of constraints. It is very convenient to define a Flex-Algorithm that uses bandwidth based metric-type which can be used for carrying high bandwidth traffic. In this regard, it's useful to define additional metric-type and additional bandwidth related constraints to simplify the operations.

This document specifies a new metric-type to be used in Flex-Algorithm described in section [Section 2](#). Additional Flexible Algorithm Definition (FAD) constraints are defined in section [Section 3](#). [Section 4](#) defines mechanisms to automatically calculate metric based on parameters defined in a FAD constraint and the advertised Maximum Link Bandwidth of the link.

[2.](#) Bandwidth Metric Advertisement

ISIS and OSPF advertise link metric in their respective link information. Multiple types of metric are supported, IGP cost, te-metric defined in [[RFC5305](#)] and [[RFC3630](#)] and delay metric defined in [[RFC8570](#)] and [[RFC7471](#)]. A brownfield network might have deployed legacy transport mechanisms using igp-cost and te-metric which continue to run during migration period. In this brownfield network if the operator wants to introduce two Flex-Algorithms, one for delay metric and another for bandwidth metric, a new metric-type to carry

The Bandwidth Metric sub-TLV MUST be advertised only once. If there are multiple Bandwidth Metric sub-TLVs advertised for a link in one or more received LSPDUs, the first one MUST be used and the subsequent ones MUST be ignored.

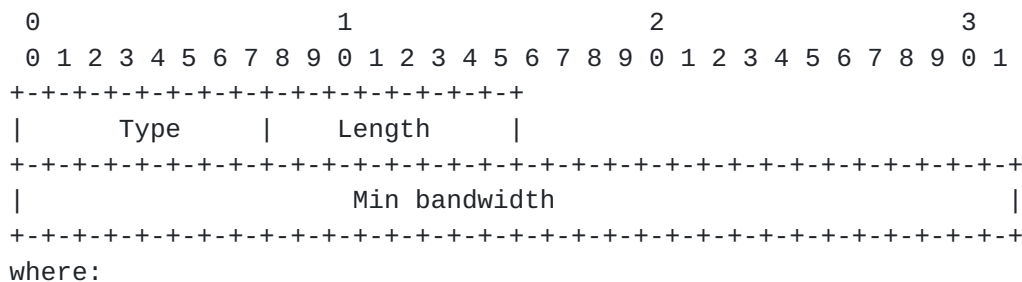
minimum link bandwidth required may change based on the applications that use the high bandwidth Flex-Algo. This document proposes a new Exclude Minimum Bandwidth constraint. When this constraint is advertised in a FAD, based on the advertised link bandwidth, the link will be pruned from the Flex-Algorithm topology if the link's advertised Maximum Link Bandwidth is below the FAD advertised Minimum bandwidth value.

Similarly, exclude maximum link delay constraint is also defined in this document. Links may have the link delay measured dynamically and advertised in delay metric in IGP. For use cases that deploy dynamic link delay measurement, A Flex-Algorithm may want to exclude links that have link delay more than a defined threshold.

3.1. ISIS FAD constraint sub-TLVs

3.1.1. ISIS Exclude Minimum Bandwidth sub-TLV

ISIS Flex-Algorithm Exclude Minimum Bandwidth sub-TLV (FAEMB) is a sub-TLV of the ISIS FAD sub-TLV. It has the following format.



where:

Type: 1

Length: 4 octets.

Min Bandwidth: link bandwidth is encoded in 32 bits in IEEE floating point format. The units are bytes per second.

Figure 3: ISIS FAEMB sub-TLV

The FAEMB sub-TLV MUST appear only once in the FAD sub-TLV. If it appears more than once, the ISIS FAD Sub-TLV MUST be ignored by the receiver. The Minimum bandwidth advertised in FAEMB sub-TLV MUST be compared with Maximum Link Bandwidth advertised in sub-sub-TLV 9 of ASLA sub-TLV [RFC 8919]. If L-Flag is set in the ASLA sub-TLV, the Minimum bandwidth advertised in FAEMB sub-TLV MUST be compared with Maximum Link Bandwidth as advertised by the sub-TLV 9 of the TLV 22/222/23/223/141 [RFC 5305] as defined in [RFC8919] Section 4.2. If the Maximum Link Bandwidth is lower than the Minimum link bandwidth

advertised in FAEMB sub-TLV, the link MUST be excluded from the Flex-Algorithm topology. If a link does not have the Maximum Link Bandwidth advertised but the FAD contains this sub-TLV, then that link MUST be included in the topology and proceed to apply further pruning rules for the link.

3.1.2. ISIS Exclude Maximum Delay sub-TLV

ISIS Flex-Algorithm Exclude Maximum Delay sub-TLV (FAEMD) 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      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     max link delay                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
where:
```

Type: TBD

Length: 3 octets

Max link delay: Maximum link delay in microseconds

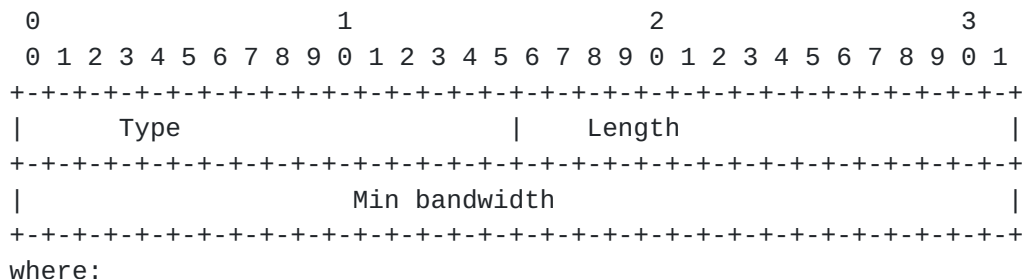
Figure 4: ISIS FAEMD sub-TLV

The FAEMD sub-TLV MUST appear only once in the FAD sub-TLV. If it appears more than once, the ISIS FAD Sub-TLV MUST be ignored by the receiver. The Maximum link delay advertised in FAEMD sub-TLV MUST be compared with Min Unidirectional Link Delay advertised in sub-sub-TLV 34 of ASLA sub-TLV [RFC 8919]. If L-Flag is set in the ASLA sub-TLV, the Maximum link delay advertised in FAEMD sub-TLV MUST be compared with Min Unidirectional Link Delay as advertised by the sub-TLV 34 of the TLV 22/222/23/223/141 [RFC 8570] as defined in [RFC8919] Section 4.2. If the Min Unidirectional Link Delay value is higher than the Maximum link delay advertised in FAEMD sub-TLV, the link MUST be excluded from the Flex-Algorithm topology. If a link does not have the Min Unidirectional Link Delay advertised but the FAD contains this sub-TLV, then that link MUST be included in the topology and proceed to apply further pruning rules for the link.

3.2. OSPF FAD constraint sub-TLVs

3.2.1. OSPF Exclude Minimum Bandwidth sub-TLV

OSPF Flex-Algorithm Bandwidth Exclusion sub-TLV (FAEMB) is a sub-TLV of the OSPF FAD TLV. It has the following format.



Type: TBD

Length: 4 octets.

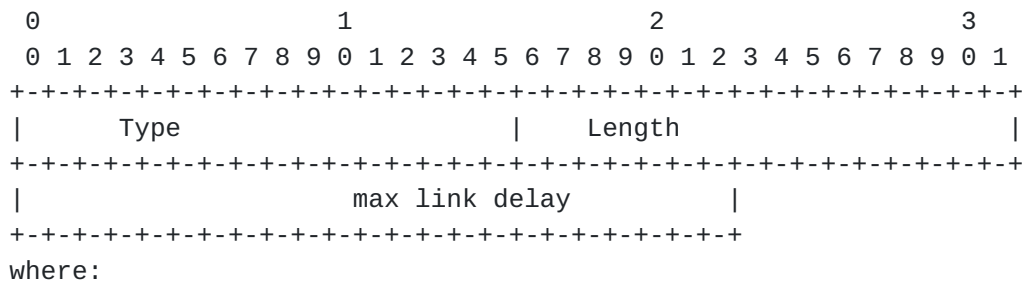
Min Bandwidth: link bandwidth is encoded in 32 bits in IEEE floating point format. The units are bytes per second.

Figure 5: OSPF FAEMB sub-TLV

The FAEMB sub-TLV MUST appear only once in the FAD sub-TLV. If it appears more than once, the OSPF FAD TLV MUST be ignored by the receiver. The Maximum Link Bandwidth as advertised by the sub-sub-TLV 23 of ASLA [[RFC 8920](#)] MUST be compared against the Minimum bandwidth advertised in FAEMB sub-TLV. If the link bandwidth is lower than the Minimum bandwidth advertised in FAEMB sub-TLV, the link MUST be excluded from the Flex-Algorithm topology. If a link does not have the Maximum Link Bandwidth advertised but the FAD contains this sub-TLV, then that link MUST be included in the topology and proceed to apply further pruning rules for the link.

3.2.2. OSPF Exclude Maximum Delay sub-TLV

OSPF Flex-Algorithm Exclude Maximum Delay sub-TLV (FAEMD) is a sub-TLV of the OSPF FAD TLV. It has the following format.



Type: TBD

Length: 3 octets

Max link delay: Maximum link delay in microseconds

Figure 6: OSPF FAEMD sub-TLV

The FAEMD sub-TLV MUST appear only once in the OSPF FAD TLV. If it appears more than once, the OSPF FAD TLV MUST be ignored by the receiver. The Min Unidirectional Link Delay as advertised by sub-sub-TLV 12 of ASLA sub-TLV [[RFC 8920](#)], MUST be compared against the Maximum delay advertised in FAEMD sub-TLV. If the Min Unidirectional Link Delay is higher than the Maximum delay advertised in FAEMD sub-TLV, the link MUST be excluded from the Flex-Algorithm topology. If a link does not have the Min Unidirectional Link Delay advertised but the FAD contains this sub-TLV, then that link MUST be included in the topology and proceed to apply further pruning rules for the link.

4. Automatic Metric Calculation

Networks which are designed to be highly regular and follow uniform metric assignment may want to further simplify the operations by automatically calculating the metric based on a reference bandwidth or a staircase metric assignment based on bandwidth thresholds. When a FAD advertises metric-type as Bandwidth Metric and the link does not have the Bandwidth Metric advertised, automatic metric derivation can be used with additional FAD constraint advertisements as described in this section. Based on the advertised rules, every node automatically calculates the link metric of all the links in the Flex-Algorithm before running SPF algorithm for the Flex-Algorithm.

Based on the delay in learning the link bandwidth changes on a remote link, there may be possibility of micro-loops which is no different from IGP susceptibility to micro-loops during metric change. The micro-loop avoidance procedures described in [\[I-D.bashandy-rtgwg-segment-routing-uloop\]](#) can be used to avoid micro-loops when the automatic metric calculation is deployed.

4.1. Automatic Metric Calculation Modes

4.1.1. Simple Mode

In simple mode, the Maximum Link Bandwidth of a single Layer 3 link is used to derive the metric. This mode is more suitable for deployments that use Layer 2 bundles for parallel links between two nodes. This layer 3 link may consist of Layer 2 bundles. In such cases, the bandwidth of the link may vary based on Layer 2 links going up/down.

4.1.2. Interface Group Mode

Simple mode of metric calculation may not work well when there are multiple parallel layer 3 interfaces between two nodes.




Figure 7: Parallel interfaces

In the above diagram, there are two parallel links between B->C, C->F, F->D. Let us assume the link bandwidth is uniform 10Gbps on all links. When Simple mode of metric derivation is used, the metric is derived as 10 on all links. Traffic will be forwarded B->E->D. Since the bandwidth is higher B->C->F->D path, the requirement is to be able to assign smaller metric based on cumulative bandwidth of the parallel links. Interface group mode is suitable for deployments that do not use Layer 2 bundles.

In the interface group mode, every node MUST identify the set of parallel links between a pair of nodes based on IGP link advertisements and MUST consider cumulative bandwidth of the parallel links while arriving at the metric of each link.

4.2. Automatic Metric Calculation Methods

In automatic metric calculation for simple and interface group mode, Maximum Link Bandwidth of the links is used to derive the metric. There are two types of automatic metric derivation methods.

1. Reference bandwidth method
2. Bandwidth thresholds method

4.2.1. Reference Bandwidth method

In many networks, the metric is assigned relative to the link bandwidth. The metric is derived by dividing the reference bandwidth by the advertised Maximum Link Bandwidth. It is useful to advertise reference bandwidth in FAD constraints to simplify operations. When there are changes to the reference bandwidth it can be easily changed on FAD server rather than having to change it on every node. In order to ensure that small bandwidth changes do not change the link metric, it is useful to define a round-off value. The link bandwidth will be rounded-off with this value before deriving the metric. For example,

reference bandwidth = 1000G

round off value = 20G

The derived metric is 10 for link bandwidth in the range 100G to 119G

4.2.2. Bandwidth Thresholds method

The round-off bw along with reference bandwidth approach described above provides a uniform metric value for a range of link bandwidth. In certain cases there may be a need to define non-proportional metric values for the varying ranges of link bandwidth. For Example, 10G to 30G metric value is 100, 30G to 70G metric value is 50, greater than 70G metric is 10. In order to support this, a staircase metric based on bandwidth threshold is supported in the FAD. This advertisement contains a set of threshold values and their associated metric.

4.3. ISIS FAD constraint sub-TLVs for automatic metric calculation

4.3.1. Reference Bandwidth sub-TLV

This section provides FAD constraint advertisement details for the reference bandwidth method of metric calculation as described in [Section 4.2.1](#). The Flexible Algorithm Definition Reference Bandwidth Sub-TLV (FADRB Sub-TLV) is a Sub-TLV of the ISIS FAD sub-TLV. It has the following format:


```

  0 1 2 3 4 5 6 7
+-+--+--+--+--+
|G| | |           |
+-+--+--+--+--+

```


G-flag: when set, interface group Mode MUST be used to derive total link bandwidth.

Staircase bandwidth threshold and associated metric values.

Bandwidth Threshold 1 Min.: Minimum Link Bandwidth is encoded in 32 bits in IEEE

floating point format. The units are bytes per second.

Bandwidth Threshold 1 Max.: Maximum Link Bandwidth is encoded in 32 bits in IEEE

floating point format. The units are bytes per second.

Threshold Metric 1 : metric value range (1 - 4,261,412,864)

Figure 9: ISIS FADBT sub-TLV

When G-flag is set, the cumulative bandwidth of the parallel links is computed as described in section [Section 4.1.2](#). If G-flag is not set, the advertised Maximum Link Bandwidth is used.

When the computed link bandwidth is less than Bandwidth Threshold 1 Min, the MAX_METRIC value of 4,261,412,864 MUST be assigned as the Bandwidth Metric on the link during Flex-Algorithm SPF calculation.

When the computed link bandwidth is greater than or equal to Bandwidth Threshold 1 Min AND less than Bandwidth Threshold 1 Max, Threshold Metric 1 MUST be assigned as the Bandwidth Metric on the link during Flex-Algorithm SPF calculation.

Similarly, when the computed link bandwidth is greater than or equal to Bandwidth Threshold 1 Max AND less than Bandwidth Threshold 2 Max, Threshold Metric 2 MUST be assigned as the Bandwidth Metric on the link during Flex-Algorithm SPF calculation.

Similarly, when the computed link bandwidth is greater than or equal to Bandwidth Threshold n-1 Max, Threshold Metric n MUST be assigned as the Bandwidth Metric on the link during Flex-Algorithm SPF calculation.

The ISIS FADBT Sub-TLV MUST 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.

A FAD MUST NOT contain both FADBT sub-TLV and FADRB sub-TLV. If both these sub-TLVs are advertised in the same FAD for a Flexible Algorithm, the FAD must be ignored by the receiver.

If a Bandwidth Metric sub-TLV is advertised for a link, the Flex-

Algorithm calculation MUST use the Bandwidth Metric advertised on the link, and MUST NOT use the automatically derived metric for that link.

4.4. OSPF FAD constraint sub-TLVs for automatic metric calculation

4.4.1. Reference Bandwidth sub-TLV

The Flexible Algorithm Definition Reference Bandwidth Sub-TLV (FADRB Sub-TLV) is a Sub-TLV of the OSPF FAD TLV. It has the following format:

Round-off Bandwidth value is used to make sure the metric does not change when there is smaller change in the link bandwidth. The OSPF FADRB Sub-TLV MUST 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. If a Bandwidth Metric sub-TLV is advertised for a link, the Flex-Algorithm calculation MUST use the Bandwidth Metric

advertised on the link, and MUST NOT use the automatically derived metric for that link.

4.4.2. Bandwidth Threshold sub-TLV

The Flexible Algorithm Definition Bandwidth Thresholds Sub-TLV (FADBT Sub-TLV) is a Sub-TLV of the OSPF FAD 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      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Reserved  |G|  Flags  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Bandwidth Threshold 1 Min.      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Threshold Metric 1      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Bandwidth Threshold 1 Max.      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Threshold Metric 2      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Bandwidth Threshold 2 Max.      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Threshold Metric n-1      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Bandwidth Threshold n-1 Max.      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Threshold Metric n      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

where:

Type: TBD

Length: $2 + n \times 8$ octets. Here n is equal to number of Threshold Metrics specified.

n MUST be greater than or equal to 2.

Flags:


```

      0 1 2 3 4 5 6 7
    +-+-+-+-+-+-+-+-+
    |G| | |         |
    +-+-+-+-+-+-+-+-+

```

G-flag: when set, interface group Mode MUST be used to derive total link bandwidth.

Staircase bandwidth threshold and associated metric values.

Bandwidth Threshold 1 Min.: Minimum Link Bandwidth is encoded in 32 bits in IEEE

floating point format. The units are bytes per second.

Bandwidth Threshold 1 Max.: Maximum Link Bandwidth is encoded in 32 bits in IEEE

floating point format. The units are bytes per second.

Threshold Metric 1 : metric value range (1 - 4,294,967,296)

Figure 11: OSPF FADBT sub-TLV

When G-flag is set, the cumulative bandwidth of the parallel links is computed as described in section [Section 4.1.2](#). If G-flag is not set, the advertised Maximum Link Bandwidth is used.

When the computed link bandwidth is less than Bandwidth Threshold 1 Min, the MAX_METRIC value of 4,294,967,296 MUST be assigned as the Bandwidth Metric on the link during Flex-Algorithm SPF calculation.

When the computed link bandwidth is greater than or equal to Bandwidth Threshold 1 Min AND less than Bandwidth Threshold 1 Max, Threshold Metric 1 MUST be assigned as the Bandwidth Metric on the link during Flex-Algorithm SPF calculation.

Similarly, when the computed link bandwidth is greater than or equal to Bandwidth Threshold 1 Max AND less than Bandwidth Threshold 2 Max, Threshold Metric 2 MUST be assigned as the Bandwidth Metric on the link during Flex-Algorithm SPF calculation.

Similarly, when the computed link bandwidth is greater than or equal to Bandwidth Threshold n-1 Max, Threshold Metric n MUST be assigned as the Bandwidth Metric on the link during Flex-Algorithm SPF calculation.

The OSPF FADBT Sub-TLV MUST 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.

A FAD MUST NOT contain both FADBT sub-TLV and FADRB sub-TLV. If both these sub-TLVs are advertised in the same FAD (winner FAD) for a Flexible Algorithm, the receiving nodes MUST stop participating in such Flexible-Algorithm

If a Bandwidth Metric sub-TLV is advertised for a link, the Flex-Algorithm calculation MUST use the Bandwidth Metric advertised on the link, and MUST NOT use the automatically derived metric for that link.

5. Bandwidth metric considerations

This section specifies the rules of deriving the Bandwidth Metric if and only if the winning FAD for the Flex-Algorithm specifies the metric-type as "Bandwidth Metric".

1. If the the Bandwidth Metric sub-TLV is advertised for the link as described in [Section 2](#), it MUST be used during the Flex-Algorithm calculation.
2. If the Bandwidth Metric sub-TLV is not advertised for the link and the winning FAD for the Flex-Algorithm does not specify the automatic bandwidth metric calculation (as defined in [Section 4](#)), the Bandwidth Metric is considered as not being advertised for the link.
3. If the Bandwidth Metric sub-TLV is not advertised for the link and the winning FAD for the Flex-Algorithm specifies the automatic bandwidth metric calculation (as defined in [Section 4](#)), the Bandwidth Metric metric MUST be automatically calculated as per the procedures defined in [Section 4](#). If the Bandwidth Metric can not be calculated due to lack of Flex-Algorithm specific ASLA advertisement of sub-sub-TLV 9 [[RFC 8919](#)], or in case of IS-IS, in presence of the L-Flag in the Flex-Algorithm specific ASLA advertisement the lack of sub-TLV 9 in the TLV 22/222/23/223/141 [[RFC 5305](#)], the Bandwidth Metric is considered as not being advertised for the link.

6. Calculation of Flex-Algorithm paths

Two new additional rules are added to the existing rules in the Flex-rules specified in sec 13 of [[I-D.ietf-lsr-flex-algo](#)].

6. Check if any exclude FAEMB rule is part of the Flex-Algorithm definition. If such exclude rule exists and the link has Maximum Link Bandwidth advertised, check if the link bandwidth satisfies the FAEMB rule. If the link does not satisfy the FAEMB rule, the link MUST be pruned from the computation.
7. Check if any exclude FAEMD rule is part of the Flex-Algorithm definition. If such exclude rule exists and the link has Min Unidirectional link delay advertised, check if the link delay

satisfies the FAEMD rule. If the link does not satisfy the FAEMD rule, the link MUST be pruned from the computation.

7. Backward Compatibility

8. Security Considerations

TBD

9. IANA Considerations

9.1. IGP Metric-Type Registry

Type: Suggested 3 (TBA)

Description: Bandwidth metric

Reference: This document

9.2. ISIS Sub-Sub-TLVs for Flexible Algorithm Definition Sub-TLV

Type: Suggested 6 (TBA)

Description: ISIS Exclude Minimum Bandwidth sub-TLV

Reference: This document [Section 3.1.1](#)

Type: Suggested 7 (TBA)

Description: ISIS Exclude Maximum Delay sub-TLV

Reference: This document [Section 3.1.2](#)

Type: Suggested 8 (TBA)

Description: ISIS Reference Bandwidth sub-TLV

Reference: This document [Section 4.3.1](#)

Type: Suggested 9 (TBA)

Description: ISIS Threshold Metric sub-TLV

Reference: This document [Section 4.3.2](#)

9.3. OSPF Sub-TLVs for Flexible Algorithm Definition Sub-TLV

Type: Suggested 6 (TBA)

Description: OSPF Exclude Minimum Bandwidth sub-TLV

Reference: This document [Section 3.2.1](#)

Type: Suggested 7 (TBA)

Description: OSPF Exclude Maximum Delay sub-TLV

Reference: This document [Section 3.2.2](#)

Type: Suggested 8 (TBA)

Description: OSPF Reference Bandwidth sub-TLV

Reference: This document [Section 4.4.1](#)

Type: Suggested 9 (TBA)

Description: OSPF Threshold Metric sub-TLV

Reference: This document [Section 4.4.2](#)

9.4. Sub-TLVs for TLVs 22, 23, 25, 141, 222, and 223

Type: Suggested 45 (TBA)

Description: Bandwidth metric

Reference: This document [Section 2.1](#)

9.5. Sub-sub-TLV Codepoints for Application-Specific Link Attributes

Type: Suggested 45 (TBA)

Description: Bandwidth metric

Reference: This document [Section 2.1](#)

9.6. OSPFv2 Extended Link TLV Sub-TLVs

Type: Suggested 45 (TBA)

Description: Bandwidth metric

Reference: This document [Section 2.2](#)

[9.7.](#) Types for sub-TLVs of TE Link TLV (Value 2)

Type: Suggested 45 (TBA)

Description: Bandwidth metric

Reference: This document [Section 2.2](#)

[10.](#) Acknowledgements

Many thanks to Chris Bowers, Krzysztof Szarcowitz, Julian Lucek, Ram Santhanakrishnan for discussions and inputs.

[11.](#) Contributors

1. Salih K A

Juniper Networks

salih@juniper.net

[12.](#) References

[12.1.](#) 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-13](#) (work in progress), October 2020.

[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>>.

[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>>.

[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>>.

- [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>>.

12.2. Informative References

- [I-D.bashandy-rtgwg-segment-routing-uloop]
Bashandy, A., Filsfils, C., Litkowski, S., Decraene, B., Francois, P., and P. Psenak, "Loop avoidance using Segment Routing", [draft-bashandy-rtgwg-segment-routing-uloop-10](#) (work in progress), December 2020.
- [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>>.
- [RFC5311] McPherson, D., Ed., Ginsberg, L., Previdi, S., and M. Shand, "Simplified Extension of Link State PDU (LSP) Space for IS-IS", [RFC 5311](#), DOI 10.17487/RFC5311, February 2009, <<https://www.rfc-editor.org/info/rfc5311>>.
- [RFC5316] Chen, M., Zhang, R., and X. Duan, "ISIS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", [RFC 5316](#), DOI 10.17487/RFC5316, December 2008, <<https://www.rfc-editor.org/info/rfc5316>>.
- [RFC7471] Giacalone, S., Ward, D., Drake, J., Atlas, A., and S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", [RFC 7471](#), DOI 10.17487/RFC7471, March 2015, <<https://www.rfc-editor.org/info/rfc7471>>.
- [RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", [RFC 8570](#), DOI 10.17487/RFC8570, March 2019, <<https://www.rfc-editor.org/info/rfc8570>>.

Authors' Addresses

Shraddha Hegde
Juniper Networks Inc.
Exora Business Park
Bangalore, KA 560103
India

Email: shraddha@juniper.net

William Britto
Juniper Networks Inc.

Email: bwilliam@juniper.net

Rajesh Shetty
Juniper Networks Inc.

Email: mrajesh@juniper.net

Bruno Decraene
Orange

Email: bruno.decraene@orange.com

Peter Psenak
Cisco Systems

Email: ppsenak@cisco.com

Tony Li
Arista Networks

Email: tony.li@tony.li

