Workgroup: IDR
Internet-Draft:
draft-ssangli-idr-bgp-generic-metric-aigp-07
Published: 9 November 2023
Intended Status: Standards Track
Expires: 12 May 2024
Authors: S. Sangli, Juniper Networks Inc.
S. Hegde, Juniper Networks Inc.
R. Das, Juniper Networks Inc.
B. Decraene, Orange
B. Wen, Comcast
M. Kozak, Comcast
J. Dong, Huawei
L. Jalil, Verizon
K. Talaulikar, Cisco

Generic Metric for the AIGP attribute

Abstract

This document defines extensions to the AIGP attribute to carry Generic Metric sub-types. This is applicable when multiple domains exchange BGP routing information. The extension will aid in intent-based end-to-end path selection.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Large Networks belonging to an enterprise may consist of nodes in the order of thousands and may span across multiple IGP domains where each domain can run separate IGPs or levels/areas. BGP may be used to interconnect such IGP domains, with one or more IGP domains within an Autonomous System. The enterprise network can have multiple Autonomous Systems and BGP may be employed to provide connectivity between these domains. Furthermore, BGP can be used to provide routing over a large number of such independent administrative domains.

The traffic types have evolved over years and operators have resorted to defining different metric types within a IGP domain (ISIS or OSPF) for IGP path computation. An operator may want to create an end-to-end path that satisfy certain intent. The intent could be to create end-to-end path that minimizes one of the metric-types. Some metrics can be assigned administratively by an operator and they are described in the base ISIS, OSPF specifications. Other metrics, for example, are the Traffic Engineering Default Metric defined in [RFC5305] and [RFC3630] , Min Unidirectional delay metric defined in [RFC8570] and [RFC7471] . There may be other metrics such as jitter, reliability, fiscal cost, etc. that an operator may wish to express as the cost of a link. The procedures mentioned in the above specifications describe the IGP path computation within IGP domains.

With the advent of 5G applications and Network Slicing applications, an operator may wish to provision end-to-end paths across multiple domains to cater to traffic constraints. This is also known as intent-based inter-domain routing. The problem space and requirements are described in [I-D.draft-hr-spring-intentaware-routing-using-color]

The Clasful Transport Planes as described in [I-D.draft-ietf-idr-bgp-ct] and and Color-Based Routing as described in [I-D.draft-ietf-idr-bgp-car] describe how end-to-end intent-based paths can be established. The proposal described in this document can be used in conjunction with such architectures.

When multiple domains are interconnected via BGP, protocol extensions for advertising best-external path and/or ADDPATH as described in [RFC7911] are employed to take advantage of network connectivity thus providing alternate paths. The Color-Based Routing and Classful Transport Planes routing proposals describe approaches that result in alternate paths for a reaching one destination. During the BGP best path computation, the step(e) as per section 9.1.2.2 of [RFC4271] , the interior cost of a route as determined via the IGP metric value can be used to break the tie. In a network spanning multiple IGP domains, the AIGP TLV encoded within the AIGP attribute described in [RFC7311] can be used to compute the AIGP-enhanced interior cost to be used in the decision process for selecting the best path as documented in section 2 of [RFC7311] . The [RFC7311] specifies how AIGP TLV can carry the accumulated IGP metric value.

There is a need to synchronize the metric-type values carried between IGP and BGP in order to avoid operational overhead of translation between them. The existing AIGP TLV carries a TLV type and metric-value where TLV type does not map to IGP metric-types defined in the IGP metric-type registry. Hence there is a need to provide a generic metric template to embed the IGP metric-type values within the AIGP attribute. This document extends the AIGP attribute for carrying Generic-Metric TLV and the well-defined sub metric types. This document also provides procedures for handling Generic-Metric during the BGP best path computation.

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. Multiple Metric types

Consider the network as shown in Figure 1. The network has multiple domains. Each domain runs a separate IGP instance. Within each domain iBGP sessions are established between the PE routers. eBGP sessions are established between the Border Routers across domains. An operator wishes to compute end-to-end path optimized for a metric-type delay. Each domain will be enabled to compute the IGP paths based on metric-type delay. Such values should also be propagated to the adjacent domains for effective end-to-end path computation.

```
    |   IBGP    |  EBGP  |   IBGP    |  EBGP  |   IBGP    |

    +----------+        +----------+        +----------+
    |          |        |          |        |          |
    |        ASBR1+--+ASBR2      ASBR3+--+ASBR4        |
    |          |        |          |        |          |
 PE1+ Domain1  |        | Domain2  |        |  Domain3 |
    |          |        |          |        |          |
    |        ASBR5+---+ASBR6     ASBR7+--+ASBR8        |
    |          |        |          |        |          |
    +----------+        +----------+        +----------+

    |  ISIS1   |        |   ISIS2  |        |  ISIS3   |
```

Figure 1: WAN Network

The AIGP TLV in the AIGP attribute as specified in [RFC7311] supports the default IGP-metric. If all domains use default IGP-metric cost, then one can compute the end-to-end path with shortest default IGP-metric cost. However if an operator wishes to compute the end-to-end path with metric other than IGP cost, we need additional extensions to the AIGP attribute for carry the metric-types and metric values.

The [I-D.ietf-lsr-flex-algo-bw-con] proposes a generic metric type that can embed multiple metric types within it. It supports both standard metric-types and user-defined metric-types. This document leverages the generic-metric draft and proposes extensions to the AIGP attribute to carry Generic Metric TLV as specified below.

4. Issues with RFC7311

The following procedures are not clearly described in [RFC7311] .

* The section 3 describes "When an AIGP attribute is created, it SHOULD contain no more than one AIGP TLV. However, if it contains more than one AIGP TLV, only the first one is used as described in Sections 3.4 and 4. In the remainder of this document, we will use the term value of the AIGP TLV to mean the value of the first AIGP TLV in the AIGP attribute. Any other AIGP TLVs in the AIGP attribute MUST be passed along unchanged if the AIGP attribute is passed along."

* ....One MUST interpret that more than one TLV of a particular type (i.e. AIGP TLV metric-type 1) can be present in the update and only the first occurance MUST be analysed. All other TLVs (type 2 or type 3 etc.) MUST be passed along unchanged if AIGP attribute is passed along.

*The section 3.2 describes "Note that an AIGP attribute MUST NOT be considered to be malformed because it contains more than one TLV of a given type or because it contains TLVs of unknown types."

*....One MUST interpret that opaque TLVs (TLVs with type 2 or type 3 for example) MUST be passed along if ADVERTISE_AIGP_ATTRIBUTE has been enabled to a neighbor.

*Section 3.3 describes "The AIGP attribute MUST NOT be sent on any BGP session for which AIGP_SESSION is disabled."

*....While maintaining the non-transitivity is important, it is also important to provide accumulated cost end-to-end across domains. If there are more than one TLVs in the AIGP attribute, it becomes important to define the behaviour of which TLV gets updated and sent across domains.

*The rules for route redistribution is not clearly described.

*....When a BGP route is redistributed, should AIGP metric-value be used directly as the cost in IGP or should there be a policy to modify AIGP metric-value before redistributing the route into IGP. It is important to define the behaviour of route redistribution metric conversion when redistribution occurs on multiple domains along the path.

5. Generic Metric TLV

This document proposes a new TLV : Generic-Metric TLV in the AIGP attribute. This will carry the metric type and metric value used in the network. The format is shown below.

```
0                   1                   2                   3
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 2
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Type      |              Length          |  metric-type  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| metric-flags| metric-value                                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+.........................
```

Figure 2: Generic-Metric TLV

Generic-Metric TLV Type (1 octet): Code point to be assigned by IANA

Generic-Metric TLV Length (2 octets): Value 10

Generic-Metric TLV Value (10 octets): 3 sub-fields as shown below:

1. metric-type (1 octet): Value of metric-types from IGP Protocol registry.

2. metric-flags (1 octet): Bits defined below.

3. metric-value (8 octets): Value range (0 - 0xffffffffffffffff)

The metric-flags carry additional information about the Generic-Metric.

```
 7 6 5 4 3 2 1 0
+-+-+-+-+-+-+-+-+
|R|R|R|R|R|R|N|I|
+-+-+-+-+-+-+-+-+
```

Figure 3: Generic-Metric Flags

Bit I : Represents incomplete/discontinuous metric accumulation for the end-to-end path. 1 indicates discontinuous, 0 indicates continuous.

Bit N : Represents normalization. 1 indicates metric normalization has been applied. 0 indicates no normalization has been applied.

Bit R : Reserved for future use. Reset to zero by the sender and ignored by the receiver.

6. Usage of Generic-Metric TLV

1. When a BGP speaker wishes to generate AIGP attribute with Generic-Metric TLV for a prefix, it MUST perform the following procedures.

-The procedures specified in [RFC7311] section 3.4 should be followed that describes creation of attribute, modifications by the originator and non-originator of the route in addition to the following procedures.

-The domain can adopt more than one metric type to represent the intent, hence the originator BGP speaker can encode more than one Generic-Metric TLV, each TLV carrying different metric type as defined in the IGP Protocol Registry.

-The type of metric used in the local domain and as specified in the IGP Protocol registry must be encoded in the metric-type sub-field. The value of the metric or cost to reach the prefix being advertised must be encoded in the metric-value sub-field, normalized if required. This is the cost or the distance to the destination prefix from the advertising BGP speaker which sets itself as the next hop as described in section 3.4 of [RFC7311].

-Repeated metric changes may cause large number of BGP updates to get generated and be propagated throughout the network. In order to avoid that, a configurable threshold is defined. If the difference between the new metric-value and the advertised metric-value is less than the configured threshold, the update MAY be suppressed. For each of type of metric used in the domain, if the new metric-value encoded in Generic-Metric TLV is above the configured threshold, a new BGP update containing the new set of metric-values SHOULD be advertised.

- The "I" bit of the metric-flags MUST be reset to zero if the BGP speaker is the originator of the AIGP attribute. If the IGP cost to reach the next hop is normalized to the type of the metric in the metric-type sub-field, the "N" bit of the metric-flags sub-field MUST be set to 1, else it MUST be reset to zero.

- Procedures for defining the cost to reach a next hop for various metric-types is outside the scope of this document.

2. When a BGP speaker wishes to send a BGP update attaching the AIGP attribute, it must validate if that session has been enabled for sending the AIGP attribute as per procedures mentioned in [RFC7311] .

3. When a BGP speaker receives a BGP update that has a route to T with next hop N and has the AIGP attribute with Generic-Metric TLV it MUST perform the following procedures.

- It must validate if that session has been enabled to receive the AIGP attribute as per rules mentioned in [RFC7311] .

- There can be more than one Generic-Metric TLV, each carrying different metric types. The BGP speaker must process every Generic-Metric TLV.

- For each of the Generic-Metric TLVs present in the AIGP attribute, if the BGP speaker recognizes the type of the metric encoded in the metric-type sub-field, it must process the metric-value and metric-field sub-fields of the Generic-Metric TLV.

- If the BGP speaker does not recognize the type of metric encoded in metric-type subfield of the TLV, then it must set the "I" bit in the metric-flags to 1 before propagating to other BGP speakers and must continue to process the next Generic-Metric TLV if present. If the BGP speaker does not recognize any metric-type in the Generic-Metric TLVs, it must follow the BGP decision procedure as specified in [RFC7311].

- If the type of the metric for resolving the next hop N matches with the metric-type of Generic-Metric TLV of the AIGP attribute, then the metric-value sub-field must be used in the

AIGP-enhanced interior cost computation as specified in the next section.

- If the metric-type of the path used for resolving the next hop N does not match with the metric-type of Generic-Metric TLV of the AIGP attribute, then the BGP speaker may normalize the cost of the path used for resolving the next hop before using it in the AIGP-enhanced cost computation. A policy may be used to provide the metric normalization. Additionally, the BGP speaker must set the "N" bit to indicate that metric normalization has been done before propagating the Generic-Metric TLV to other BGP speakers.

- If the BGP speaker modifies the next hop it must update the Generic-Metric TLV(s).

7. Updates to Decision Procedure

This section follows the approach as laid out in [RFC7311] to select the best path when the route has AIGP attribute with Generic-Metric TLV. The domain that the router R belongs to, may support different intent based paths represented via different types of metric. The following describes procedures in addition to the general procedure described in section 4 of [RFC7311] .

When R receives a route T with next hop N and the AIGP attribute with one or more Generic-Metric TLVs, for each Generic-Metric TLV the BGP speaker MUST perform following procedures.

If the metric-type sub-field matches with the type of the metric for the path used for resolving the next hop N, the AIGP-enhanced interior cost should be computed as below.

Let m be the cost to reach the next hop N that IGP uses for its path computation as described in [RFC7311] .

If the type of the metric for the path used for resolving the next hop N does not match with the metric-type sub-field of the Generic-Metric TLV, the cost of the path to reach next hop N may be normalized. The normalized metric value can be zero, maximum metric value or scaled up (multiple of a positive number).

Let m be the normalized value of the cost to reach the next hop N that IGP uses for its path computation as described in [RFC7311] .

The AIGP-enhanced interior cost computation as described below will be used in the decision process as described in [RFC7311] .

Let A be the value of the value of the metric-value sub-field of the Generic-Metric TLV.

   The AIGP-enhanced interior cost will be A+m as described in
   [RFC7311] .

A path with Generic-Metric TLV and a path with AIGP TLV cannot be compared. To enable end-to-end path selection based on intent, the path with Generic-Metric TLV MUST be chosen over path with AIGP TLV. The implementation should allow a local policy to specify the preference.

A path with Generic-Metric TLV of metric-type 'a' cannot be compared with a path with Generic-Metric TLV of metric-type 'b'. The path with lower metric-type MAY be chosen as best between two paths with Generic-Metric TLV and implemented consistently across AIGP domain.

8. Use-case: Different Metrics across Domains

```
                              +--------------+
                              |   Domain2    |
                              |              |
                        ......+ASBR21  ASBR22+.....
                        .     |              |    .
        +------------+  .     |  IGP-metric  |    .  +--------------+
        |   Domain1  |  .     +--------------+    .  |    Domain4   |
        |            |  .                         .  |              |
        |      ASBR11+...                         ...+ASBR41        |
        +PE1         |                               |           PE2+
        |      ASBR12+...                         ...+ASBR42        |
        |            |  .                         .  |              |
        | IGP-metric |  .                         .  | delay-metric |
        +------------+  .                         .  +--------------+
                        .     +--------------+    .
                        .     |    Domain3   |    .
                        .     |              |    .
                        ......+ASBR31  ASBR32+.....
                              |              |
                              | delay-metric |
                              +--------------+
```

Figure 4: Different metric across network

Each domain is a separate Autonomous System. Within each domain, ASBR and PE form iBGP peering and they may employ Route Reflectors. The IGP within each domain uses domain specific metric. Domain3 and Domain4 use delay as the metric while Domain1 and Domain2 use default IGP-metric cost. ASBRs across domains form eBGP peering.

Scenario 1: Find delay-based end-to-end path from Domain1 to Domain4.

   This can be achieved by the advertising router to add the AIGP
   attribute with metric type 1 that represents delay metric. In the
   above network diagram, ASBR41 (and ASBR42) will advertise prefix

PE2-loopback with Generic-Metric TLV with delay as metric-type. The metric-value sub-field of the Generic-Metric TLV will represent the cost to reach PE2's loopback end-point from the advertising router as they will do next hop self.

In Domain3, when ASRB32 advertises the prefix PE2-loopback within the local domain, it may add cost to the metric-value, the value representing the delay introduced by the DMZ link between ASRB32 to ASBR42. When ASRBR31 advertises the prefix PE2-lookback, it will perform the following procedures.

1. Compute the delay d of the path to reach ASBR32 from which it has chosen the best path.

2. Add the above d value to the metric-value sub-field of the Generic-Metric TLV.

In Domain2 however, the local metric type is default IGP-metric. The ASBR22 may follow the procedure similar to ASBR32 and add the delay value corresponding to the DMZ link between ASBR22 and ASBR41 before advertising the path internally in Domain2. When ASBR21 computes the AIGP-enhanced interior cost, as mentioned before, it may normalize the igp cost to reach ASBR22 and may add the normalized value to the delay-metric. The ASBR21 will also update metric-flags sub-field to indicate that metric value has been normalized. In the above network example, the delay cost from ASBR21 to ASBR22 is negligible and hence delay-metric value will be unchanged.

The procedures for AIGP-enhanced interior cost computation at ASBR11 (and ASBR12) will follow DMZ delay computation procedure described above. PE1 will have two paths to reach PE2-loopback: P1 via ASBR11 (and domain2) and P2 via ASBR12 (and domain3), each having respective AIGP-enhanced interior cost representing end-to-end delay. The local metric type is default IGP-metric and hence PE1 may normalize the internal igp cost for the AIGP-enhanced interior cost computation. The BGP decision process described in Section 7 will result in delay optimized end-to-end path for PE2-loopback on PE1 that can be used to resolve the service prefixes.

Scenario 2: Provide best-effort or default IGP-metric based end-to-end path while leveraging the domain-specific delay-based metric for intra-domain path selection.

All the ASBR routers will update the Generic Metric TLV for the default IGP-metric metric-type, accumulating the cost for end-to-end path. PE1 router will have two paths (from ASBR11 and ASBR12) decorated with different best-effort default IGP-metric cost. The intra-domain path to reach the domain exit can be based on domain-specific metric-type. For example, in Domain3, ASBR31 can select lowest delay path to reach ASBR32. The ASBR and the PE routers may be configured to prefer one metric-type for end-to-end path while another metric-type for intra-domain and such configuration mechanism is outside the scope of this document.

Scenario 3: Path selection when a router along the path does not support the new type of metric.

The Domain2 implements only default IGP-metric and does not support delay-metric. When ASBR21 receives the route with AIGP attribute and the Generic-Metric TLV, the metric type delay-metric is unrecognized. The ASBR21 will update the metric-flags, setting the "I" bit to 1 indicating that accumulation is incomplete. When such a route reaches PE1, the PE1 router will have two paths, one via ASBR11 with "I" bit set and another path from ASBR12 with "I" bit reset to zero. The local policy on PE1 can provide guidance on the preference between these two paths.

9. Deployment Considerations

It can be noted that a domain may normalize the metric-value of the metric-type of the path used to resolve next hop to the metric-type present in the Generic-Metric TLV. The idea is to propagate the cost of reaching the prefix through the domain while maintaining the metric-type chosen by the originating router and domain thereby providing an end-to-end path for the desired intent. The normalization of metric types to the one carried in the AIGP attribute can be done via policy. Definition of such policies and how they can be enforced is outside the scope of this document. In topologies where there is a common router between adjacent domains that do iBGP peering, the Border router can provide the normalization.

It is important to maintain the property of IGP cost to a destination decrease as one gets closer to the destination. The AIGP-enhanced interior cost should not be allowed to decrease through the metric normalization. When adjacent domains use different metric types, the ASBR that connects two domains is better suited to pass on the metric values by setting itself as next hop.

All routers of a domain MUST compute the AIGP-enhanced interior cost as described above to be used during decision process. Within a domain, if one router R1 applies AIGP-enhanced interior cost while R2 does not, it may lead to routing loop unless some sort of tunnelling technology viz MPLS, SRv6, IP, etc. is adopted to reach the next hop. In a network where any tunnelling technology is used, one can incrementally deploy the Generic-Metric functionality. In a network without any tunnelling technology, it is recommended that all routers MUST support Generic-Metric based AIGP-enhanced interior cost computation.

In certain networks, routes may be redistributed between BGP and IGP, usually controlled via a policy. When a route is propagated across domains, a router should use AIGP metric-value of Generic-Metric TLV, optionally modified via the local policy as the IGP cost during route redistribution in to IGP. The local policy should apply metric normalization or translation based on metric-type of Generic-Metric TLV and the metric-type adopted in the IGP.

10. Contiguity Compliance

AIGP attribute is optional and non-transitive, however new TLV might not be interpreted and/or updated by routers along the path. The contiguity of the AIGP domain across multiple IGP or AS domains is important to maintain end-to-end path of a certain intent. All the BGP routers along the path that modify the next hop should accumulate the cost and propagate the accumualated cost in the AIGP attribute. For calculating the end-to-end path for an intent expressed via a type of metric, all such routers MUST support the Generic-Metric handling for that type of metric and intent. This will assure the correct end-to-end path for the intent and the metric.

> If a router along the path did not recognize a certain type of metric present in the Generic-Metric TLV, from the "I" bit of the metric-flags, the receiving router can infer that metric accumulation is not complete and appropriate decision can be taken during the best path computation.
>
> If a router along the path did not support Generic-Metric TLV and yet propagated the AIGP attribute, the metric-flags would not indicate the discontiguity. It is recommended that operators identify such routers and upgrade them to support Generic-Metric TLV and it would bring in determinism.
>
> If a router along the path did not support Generic-Metric TLV and chose to drop the AIGP attribute, the receiving router will not be able to compute end-to-end path for the desired intent and metric type. Identifying such routers and upgrading them to support Generic-Metric TLV would deliver the desired results.

11. Backward Compatibility

When a BGP speaker receives an update with the AIGP attribute it may have Generic-Metric TLV. If the BGP speaker understands the AIGP attribute but does not understand the Generic-Metric TLV, it will process the AIGP attribute as per [RFC7311] . However when it needs to advertise the prefix to its peers it will pass on the AIGP attribute with all the TLVs including the unknown Generic-Metric TLV as per [RFC7311] . If a BGP speaker does not understand the Generic-Metric TLV, it may chose sub-optimal BGP path.

12. Security Considerations

This document does not introduce any new security considerations beyond those already specified in [RFC4271], [RFC7311] .

13. IANA Considerations

IANA is requested to assign a code point for Generic Metric TLV. The metric-type field refers to the IGP metric-type registry defined in [I-D.ietf-lsr-flex-algo-bw-con]

14. Acknowledgements

The authors would like to thank John Scudder, Jeff Haas, Robert Raszuk, Kaliraj Vairavakkalai, and Peng Shaofu for careful review and suggestions.

15. References

15.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[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>.

15.2. Informative References

[I-D.hr-spring-intentaware-routing-using-color] Hegde, S., Rao, D., Uttaro, J., Bogdanov, A., and L. Jalil, "Problem statement for Inter-domain Intent-aware Routing using Color", Work in Progress, Internet-Draft, draft-hr-spring-intentaware-routing-using-color-03, 23 October 2023, <https://datatracker.ietf.org/doc/html/draft-hr-spring-intentaware-routing-using-color-03>.

[I-D.ietf-idr-bgp-car] Rao, D., Agrawal, S., and Co-authors, "BGP Color-Aware Routing (CAR)", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-car-03, 23 October 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-car-03>.

[I-D.ietf-idr-bgp-ct] Vairavakkalai, K. and N. Venkataraman, "BGP Classful Transport Planes", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-ct-18, 5 November 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-ct-18>.

[I-D.ietf-lsr-flex-algo-bw-con] Hegde, S., Britto, W., Shetty, R., Decraene, B., Psenak, P., and T. Li, "Flexible Algorithms: Bandwidth, Delay, Metrics and Constraints", Work in Progress, Internet-Draft, draft-ietf-lsr-flex-algo-bw-con-07, 26 September

2023, <https://datatracker.ietf.org/doc/html/draft-ietf-lsr-flex-algo-bw-con-07>.

[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>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[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>.

[RFC7311] Mohapatra, P., Fernando, R., Rosen, E., and J. Uttaro, "The Accumulated IGP Metric Attribute for BGP", RFC 7311, DOI 10.17487/RFC7311, August 2014, <https://www.rfc-editor.org/info/rfc7311>.

[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>.

[RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", RFC 7911, DOI 10.17487/RFC7911, July 2016, <https://www.rfc-editor.org/info/rfc7911>.

[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

Srihari Sangli
Juniper Networks Inc.
Exora Business Park
Bangalore, KA 560103
India

Email: ssangli@juniper.net

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

Email: shraddha@juniper.net

Reshma Das
Juniper Networks Inc.
1133 Innovation Way
Sunnyvale, CA 94089
USA

Email: dreshma@juniper.net

Bruno Decraene
Orange
France

Email: bruno.decraene@orange.com

Bin Wen
Comcast
USA

Email: bin_wen@comcast.com

Marcin Kozak
Comcast
USA

Email: marcin_kozak@comcast.com

Jie Dong
Huawei
China

Email: jie_dong@huawei.com

Luay Jalil
Verizon
USA

Email: luay.jalil@verizon.com

Ketan Talaulikar
Cisco
India

Email: ketant.ietf@gmail.com