Network Internet-Draft Intended status: Standards Track Expires: July 1, 2022 Shaofu. Peng Bin. Tan ZTE Corporation December 28, 2021

BGP Metric Credit Based Routing draft-peng-idr-bgp-metric-credit-00

Abstract

This document defines an optional metric related BGP path attribute that can be advertised with BGP intent routes, to provide more clear intent information used for the next-hop resolution.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of <u>BCP 78</u> and <u>BCP 79</u>.

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 <u>https://datatracker.ietf.org/drafts/current/</u>.

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 July 1, 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents (<u>https://trustee.ietf.org/license-info</u>) 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

$\underline{1}$. Introduction	2
<u>1.1</u> . Requirements Language	<u>3</u>
$\underline{2}$. Extensions for BGP Route Metric	<u>4</u>
2.1. Extensions for Metric Type and Metric	<u>4</u>
2.2. Extensions for Metric Credit	<u>4</u>
$\underline{3}$. Process of Received Metric Credit	<u>6</u>
<u>3.1</u> . Only Total Metric Credit Provided	<u>6</u>
<u>3.2</u> . Metric Credit Piece Provided	<u>7</u>
3.3. Select Underlay Path According to Suggested Metric Credit	7
<u>4</u> . Examples	<u>7</u>
<u>4.1</u> . Example of Intent with Average Metric Credit	<u>8</u>
<u>4.2</u> . Example of Intent with Metric Credit Piece	<u>10</u>
<u>4.3</u> . Example of Intent Shared by Multiple Path	<u>12</u>
5. IANA Considerations	<u>15</u>
<u>6</u> . Security Considerations	<u>15</u>
<u>7</u> . Acknowledgements	<u>15</u>
<u>8</u> . Normative References	<u>15</u>
Authors' Addresses	<u>16</u>

<u>1</u>. Introduction

In large-scale networks across multiple domains, BGP can be used to advertise intent route and provide end-to-end intent aware path. According to corresponding intent information, a BGP intent route selects the expected underlay path for the next hop resolution. The underlay path may be an MPLS-TE tunnel, an SR policy, an IGP Flexalgo route, or a direct link, established segment by segment based on the same intent.

There are many methods to make BGP carry intent information during route advertisement.

[I-D.zhou-idr-inter-domain-lcu] describe a simple method, termed as BGP-LU color, combined with [<u>RFC7911</u>], to create multiple BGP-LU color path to the same destination, using color as intent identifier.

[I-D.kaliraj-idr-bgp-classful-transport-planes] defines new "Classful Transport" SAFI NLRI and "Transport Class" Route Target extended community, to advertise intent based routes with related Transport Class identifier as intent identifier.

[I-D.dskc-bess-bgp-car] defines new "BGP CAR" SAFI NLRI that contains a Color field as intent identifier to advertise intent based routes.

[I-D.lp-lsr-bgp-algorithm] introduces Algorithm Extended Community to advertise BGP algorithm routes, using algorithm as intent identifier.

Once a BGP speaker received an intent route advertised from neighbor, it will check the related intent template in local. The intent template could be a Color template, a Flex-algorithm Definition, or some other modules. The intent template contains a set of constraints, such as the reserved bandwidth to be provided in the path, the boundary minimum and maximum delay, the boundary delay jitter, the boundary packet loss rates, including or excluding specific nodes or links, limiting the calculation of the path in a specific virtual network, and so on. The detailed intent information itself are not recommended to be directly carried in the advertised route.

On the ingress PE node in the network, a BGP intent route to the egress PE will be selected according to the service SLA. The intent template configured on the ingress PE node is generally consistent with the service SLA. However this does not mean that the intent template configured on the intermediate nodes should also be consistent with the service SLA. For example, the SLA of the service is to "provide a path with an upper delay boundary of 100ms from ingress PE to egress PE". In this case, the upper delay of 100ms refers to end-to-end delay constraint, not for each segment. A possible method is to include different delay constraints in the intent template configured on different BGP speakers along the path. However, this static configuration method has shortcomings, because an intent template is not necessarily bound to a specific end-to-end path and may be used for multiple paths.

It can be seen that the information contained in the intent template is not enough to represent the complete intent. This often occurs on cumulable metric attributes.

This document describes extensions to carry the metric credit information in the BGP route advertisement, as a supplement to the intent template, so that the BGP speaker receiving the BGP intent route can establish (or select) the underlay path to the downstream BGP speaker with more accurate intent indication.

<u>1.1</u>. 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 <u>BCP</u> <u>14</u> [<u>RFC2119</u>] [<u>RFC8174</u>] when, and only when, they appear in all capitals, as shown here.

2. Extensions for BGP Route Metric

This section describes some extensions related to metric, so that when BGP speaker advertise intent route to upstream neighbors, it carries optional metric type, metric, and metric credit information.

<u>2.1</u>. Extensions for Metric Type and Metric

[RFC7311] defines AIGP Attribute to represent basic IGP default metric type and its cumulative value.

[I-D.ssangli-idr-bgp-generic-metric-aigp] defines extensions to the AIGP attribute to carry Generic Metric sub-types, to meet more metric types, such as Min Unidirectional delay metric defined in [RFC8570], TE Default Metric defined in [RFC5305], Bandwidth Metric defined in [I-D.ietf-lsr-flex-algo-bw-con], etc.

This document reuse the above definitions.

2.2. Extensions for Metric Credit

A new path attribute type (TBD): METRIC-CREDIT attribute, is introduced in this document. It is an optional non-transitive attribute, which is used to specify end-to-end total metric credit, estimated BGP hop counts, and metric credit piece information.

The format of attribute value is shown below.

<pre>+</pre>	+ _	
Flags (1 octet)		< 1st Source
// Network Address of Source (variable)	//	
Total Metric Credit for Source (4 octets)	 +	
Estimated BGP Hops Count for Source (1 octet)	 +	
// Current Hop Number (1 octet)	//	
// Metric Credit Piece [Hops Count] (variable)	// +	
// For other Sources	// +	< 2nd Source

Figure 1: Metric Credit Path Attribute Format

Count of Sources (1 octet): Indicates how many sources (i.e. ingress PE) have corresponding metric credit information.

Flags (1 octet): Currently three flags are defined.

Figure 2

S-Flag: Source Address Flag. The Network Address of Source field is included when S-Flag is set, and not included when unset.

F-Flag: Address Family Flag. Indicate the address family of the Network Address of Source field. 0 for 32 bits IPv4 address, 1 for 128 bits IPv6 address.

P-Flag: Piece Flag. Indicates whether specific metric credit piece information (Current Hop Number field and Metric Credit Piece [] field) is included.

Network Address of Source (variable octets): Represents the IP address of the source. When S-flag is 0, this field does not exist; When S-flag is 1 and F-flag is 0, this field is a 32-bits IPv4 address; When S-flag is 1 and F-flag is 1, this field is a 128-bits IPv6 address. If this field does not exist, it means that the metric credit information is independent of the specific source. This generally occurs in the scenario that want to control the establishment of intent paths according to the metric credit in coarse granularity. For example, in a small administrative domain, when the egress PE advertise BGP intent routes to multiple ingress PEs, only a single unified metric credit information for all source is carried.

Total Metric Credit for Source (4 octets): Represents the total metric credit from a specific ingress PE (or any source, when the Network Address of Source field does not exist) to egress PE.

Estimated BGP Hops Count for Source (1 octet): Represents the estimated number of BGP hops from a specific ingress PE (or any source, when the Network Address of Source field does not exist) to egress PE. Only those BGP speakers who modified BGP next hop to self when receiving BGP intent route are counted.

Current Hop Number (1 octet): Represents the current index of the array Metric Credit Piece[]. Note that when P-flag is 0, the

[Page 5]

"Current Hop Number" and "Metric Credit Piece[]" fields do not exist. In the BGP intent route advertisement originated from egress PE, the initial value of Current Hop Number is 0; When the BGP intent route advertisement received by an upstream BGP speaker and the BGP speaker modifies the BGP next hop to self, it read the item from the array Metric Credit Piece[] using Current Hop Number as index, to obtain the metric credit piece that is available from the current BGP speaker to the neighbor of the downstream BGP speaker. When the BGP speaker continue to advertise the route to the upstream BGP speaker neighbor and modify BGP next hop to self, it increase the Current Hop Number by 1. Note that when using Current Hop Number as index to read items from the array Metric Credit Piece[], it must avoid array out of bounds.

Metric Credit Piece[] (variable octets): Is an array. The number of items is specified by Estimated BGP Hops Count for Source. Each item occupies 3 bytes (in microseconds) for each BGP speaker along the intent path. There are strict restrictions on the use of Metric Credit Piece[], i.e., BGP intent route must be advertised to a specific ingress PE hop by hop in strict accordance with the propagated path expected by egress PE. When an exception occurs, e.g., a BGP speaker finds that the Current Hop Number value in the received BGP intent route is greater than or equal to the Estimated BGP Hops Count for Source value, then it must stop using Metric Credit Piece[] for BGP next hop resolution.

3. Process of Received Metric Credit

When a BGP speaker receives BGP intent route with metric credit information from neighbors, it obtains the intent identififer from the route advertisement, looks up the intent template locally to obtain the detailed intent information, and selects underlay path according to the intent information.

The metric credit attribute information contained in the route advertisement can provide a suggested metric credit value for this BGP speaker to select underlay path destinated to the downstream neighbor BGP speaker.

<u>3.1</u>. Only Total Metric Credit Provided

If the metric credit attribute only contains the total metric credit related to the source, but does not contain the explicit metric credit piece information, then for each source:

Let metric_residual_value = "Total Metric Credit for Source" -"metric of AIGP Attribute"

Let average_metric_credit_value = "Total Metric Credit for Source"
/ "Estimated BGP Hops Count for Source"

The suggested metric credit for the current segment is the minimum value of metric_residual_value (note: negative values need to be excluded) and average_metric_credit_value for all sources.

3.2. Metric Credit Piece Provided

If the BGP intent route advertisement also contains explicit metric credit piece information, then:

Let explicit_metric_credit_piece_value = Metric Credit
Piece[Current Hop Number]

The suggested metric credit for the current segment is the minimum value of metric_residual_value (note: negative values need to be excluded) and explicit_metric_credit_piece_value for all sources.

It is recommonded that this option is mainly used in the case that only a single source related metric credit piece information is included in the intent route that is advertised to a single ingress PE. Multiple source related metric credit piece information may bring conflicts and cause inaccurate suggested metric credit.

To include explicit BGP speaker list in the route advertisement and specify metric credit piece for each segment, will be described in the next version.

3.3. Select Underlay Path According to Suggested Metric Credit

The purpose of suggested metric credit is to restrict that the cumulative metric of the resolved underlay path cannot exceed the expected value. However, in some cases, if the BGP speaker finds that there is no underlay path that can meet the suggested value, this constraint can be relaxed appropriately by local policy.

For the BGP intent route installed on the received BGP speaker, the metric value is updated to the metric of AIGP attribute plus the cumulative metric of the underlay path for the corresponding metric type.

<u>4</u>. Examples

Internet-Draft

4.1. Example of Intent with Average Metric Credit

As shown in Figure 3, it contains two IGP domains. BGP neighbors are established between PE1 and ABR, ABR and PE2, to advertise BGP intent routes. Egress PE2 advertise its loopback route, loopback-PE2, to ABR through BGP, and carries intent identifier and metric credit attribute.

+-----P1---++ +-----P4---++ / | \ / | \ PE1-----P2-----ABR------P5-----PE2 \ | / \ | / +-----P3----++ |<-----IGP domain 1 ---->||<-----IGP domain 2 ---->| TE-path-11: PE1-P1-ABR TE-path-12: ABR-P4-PE2 TE-path-21: PE1-P3-ABR TE-path-22: ABR-P6-PE2

Figure 3: Inter-area Intent Routing

There are two services to communicate between ingress PE1 and egress PE2. The intent of the first service is that the end-to-end total delay of the transport path used shall not exceed 10ms, for the intent of the second service, it is 100ms. Since two intent aware paths need to be instantiated between the same source/destination for two services, in general, two intent identifiers, intent-1000 and intent-2000, need to be configured on the egress PE2.

The intent-1000 template may have the following configuration:

metric-type: Unidirectional Link Delay (ms)

total-metric-credit: 10 (ms)

metric-credit enabled

The intent-2000 template may have the following configuration:

metric-type: Unidirectional Link Delay (ms)

total-metric-credit: 100 (ms)

metric-credit enabled

The above intent template are also configured in other BGP speakers (i.e., ABR and PE1). However, it should be noted that since the

BGP Metric Credit

intent template contains the metric credit enabled command, these BGP speakers, for the matched intent routes that are not originated from themselvs, will not select the underlay path to the downstream BGP speaker neighbor only according to the total metric contained in the intent template. Instead, they should obtain the suggested metric credit from the received BGP intent route, and then select the appropriate underlay path that meets the suggested metric credit.

The total metric credit included in the intent template is mainly used for egress PE2 to generate metric credit information that is encoded in the originated intent route. Other BGP speakers (non initiator) cannot directly use it to select underlay paths.

PE2 advertises two BGP intent routes, <prefix = loopback-PE2, intentid = 1000> and <prefix = loopback-PE2, intent-id = 2000>, which are advertised to ABR respectively. The BGP next hop in the advertised route is set to PE2.

```
The metric related information encoded in <prefix = loopback-PE2,
intent-id = 1000> is:
```

metric-type: Unidirectional Link Delay

metric: assume an initial value of 0

metric-credit information:

Count of Sources: 1

S-Flag = 1, F-Flag = 0, P-Flag = 0

Network Address of Source: loopback-PE1

Total Metric Credit for Source: 10

Estimated BGP Hops Count for Source: 2

Similarly, the metric related information encoded in <prefix = loopback-PE2, intent-id = 2000> is:

metric-type: Unidirectional Link Delay
metric: assume an initial value of 0
metric-credit information:
 Count of Sources: 1

S-Flag = 1, F-Flag = 0, P-Flag = 0 Network Address of Source: loopback-PE1 Total Metric Credit for Source: 100 Estimated BGP Hops Count for Source: 2

After receiving the first route advetisement, ABR knows that the suggested metric credit from itself to the downstream BGP speaker neighbor PE2 is 5ms, i.e., the average metric credit 5, the residual metric credit 10, taking the minimum one, then select an ultra-low delay path not exceeding 5ms to progress PE2, assuming TE path-12 in Figure 3.

After receiving the second route advetisement, ABR knows that the suggested metric credit from itself to the downstream BGP speaker neighbor PE2 is 50ms, i.e., the average metric credit 50, the residual metric credit 100, taking the minimum one, then select a low delay path not exceeding 50ms to progress PE2, assuming TE path-22 in Figure 3.

ABR continues to advertise the above two BGP intent routes to the upstream BGP speaker neighbor PE1, where the metric is accumulated by the selected underlay path, the BGP next hop is modified to ABR, and the metric credit information is unchanged.

After receiving the first route advetisement, PE1 knows that the suggested metric credit from itself to the downstream BGP speaker neighbor ABR is 5ms, i.e., the average metric credit 5, the residual metric credit 6, taking the minimum one, then select an ultra-low delay path to ABR, assuming TE path-11 in Figure 3.

After receiving the second route advetisement, PE1 knows that the suggested metric credit from itself to the downstream BGP speaker neighbor ABR is 50ms, i.e., the average metric credit 50, the residual metric credit 60, taking the minimum one, then select a low delay path to ABR, assuming TE path-12 in Figure 3.

4.2. Example of Intent with Metric Credit Piece

Based on the first example, assume that two IGP domains are managed by a single provider, which is familiar with the performance of the network, and the propagated path of BGP intent route is clear, then, intent template may include explicit metric credit piece information.

Thus the intent routes originated from PE2 will contain more information related with metric credit piece.

BGP Metric Credit

```
The metric related information encoded in <prefix = loopback-PE2,
intent-id = 1000> is:
  metric-type: Unidirectional Link Delay
  metric: assume an initial value of 0
  metric-credit information:
     Count of Sources: 1
     S-Flag = 1, F-Flag = 0, P-Flag = 1
     Network Address of Source: loopback-PE1
     Total Metric Credit for Source: 10
     Estimated BGP Hops Count for Source: 2
     Current Hop Number: 0
     Metric Credit Piece [2]: [0] = 4, [1] = 6
Similarly, the metric related information encoded in <prefix =
loopback-PE2, intent-id = 2000> is:
  metric-type: Unidirectional Link Delay
  metric: assume an initial value of 0
  metric-credit information:
     Count of Sources: 1
     S-Flag = 1, F-Flag = 0, P-Flag = 1
     Network Address of Source: loopback-PE1
     Total Metric Credit for Source: 100
     Estimated BGP Hops Count for Source: 2
     Current Hop Number: 0
     Metric Credit Piece [2]: [0] = 40, [1] = 60
```

After receiving the first route advetisement, ABR knows that the suggested metric credit from itself to the downstream BGP speaker

BGP Metric Credit

neighbor ABR is 4ms, i.e., the explicit metric credit piece 4, the residual metric credit 10, taking the minimum one, then select an ultra-low delay path to PE2, assuming TE path-12 in Figure 3.

After receiving the second route advetisement, ABR knows that the suggested metric credit from itself to the downstream BGP speaker neighbor ABR is 40ms, i.e., the explicit metric credit piece 40, the residual metric credit 100, taking the minimum one, then select a low delay path to PE2, assuming TE path-22 in Figure 3.

ABR continues to advertise the above two BGP intent routes to the upstream BGP speaker neighbor PE1, where the metric is accumulated by the selected underlay path, the BGP next hop is modified to ABR, and the Current Hop Number is incremented by 1.

After receiving the first route advetisement, PE1 knows that the suggested metric credit from itself to the downstream BGP speaker neighbor ABR is 6ms, i.e., the explicit metric credit piece 6, the residual metric credit 6, taking the minimum one, then select an ultra-low delay path to ABR, assuming TE path-11 in Figure 3.

After receiving the second route advetisement, PE1 knows that the suggested metric credit from itself to the downstream BGP speaker neighbor ABR is 60ms, i.e., the explicit metric credit piece 60, the residual metric credit 60, taking the minimum one, then select a low delay path to ABR, assuming TE path-12 in Figure 3.

4.3. Example of Intent Shared by Multiple Path

As shown in Figure 4, it contains three AS. BGP neighbors are established between PE1 and ASBR1, ASBR1 and ASBR2, ASBR1 and ASBR3, ASBR2 and PE2, ASBR3 and PE3, to advertise BGP intent routes carrying intent identifier and metric credit attribute.



TE-path-11:PE1-P1-ASBR1 TE-path-12:ASBR1-ASBR2 TE-path-13: ASBR2-P4-PE2 TE-path-21:PE1-P3-ASBR1 TE-path-22:ASBR1-ASBR3 TE-path-23: ASBR3-P7-PE3

Figure 4: Inter-AS Intent Routing

In this example, the same type of service needs to communicate between PE1 and PE2, PE1 and pE3. The intent of this type of service is the specific value related to the end-to-end total delay and distance of the transport path used.

PE2 and pE3 will respectively config the intent templates with the same intent identifier, e.g., intent-1000.

The intent-1000 template configured on PE2 have the following configuration:

metric-type: Unidirectional Link Delay (ms)

total-metric-credit: 200 (ms)

metric-credit enabled

The intent-1000 template also configured on PE3 have the following configuration:

metric-type: Unidirectional Link Delay (ms)

total-metric-credit: 300 (ms)

metric-credit enabled

PE2 advertises BGP intent route, <prefix = loopback-PE2, intent-id = 1000>, which is advertised to ASBR2. The metric related information encoded is:

metric-type: Unidirectional Link Delay

metric: assume an initial value of 0 metric-credit information: Count of Sources: 1 S-Flag = 1, F-Flag = 0, P-Flag = 0 Network Address of Source: loopback-PE1 Total Metric Credit for Source: 200 Estimated BGP Hops Count for Source: 3 Similarly, PE3 advertises BGP intent route, <prefix = loopback-PE3, intent-id = 1000>, which is advertised to ASBR3. The metric related information encoded is: metric-type: Unidirectional Link Delay metric: assume an initial value of 0 metric-credit information: Count of Sources: 1 S-Flag = 1, F-Flag = 0, P-Flag = 0 Network Address of Source: loopback-PE1 Total Metric Credit for Source: 300

Estimated BGP Hops Count for Source: 3

Then, ASBR2 and ASBR3 will use the suggested metric credit 66 and 100, taking minimum value from residual metric credit and average metric credit, to select the transport path for <prefix = loopback-PE2, intent-id = 1000> and <prefix = loopback-PE3, intent-id = 1000> respectively, assuming TE path-13 with cumulated metric 60 and TE path-23 with cumulated metric 100 in Figure 4.

Similarly, ASBR1 will also use suggested metric credit 66 and 100, to select the transport path for <prefix = loopback-PE2, intent-id = 1000> and <prefix = loopback-PE3, intent-id = 1000> respectively, assuming TE path-12 with cumulated metric 10 and TE path-22 with cumulated metric 10 in Figure 4.

Internet-Draft

Similarly, PE1 will also use suggested metric credit 66 and 100, to select the transport path for <prefix = loopback-PE2, intent-id = 1000> and <prefix = loopback-PE3, intent-id = 1000> respectively, assuming TE path-11 with cumulated metric 60 and TE path-21 with cumulated metric 100 in Figure 4.

<u>5</u>. IANA Considerations

This document request the METRIC_CREDIT entry to the "BGP Path Attributes" registry:

+=====+=======+======+=======+ | Value | Code | Reference | +=====+=========+ | TBD | METRIC_CREDIT | This Document | +----+

6. Security Considerations

TBD

7. Acknowledgements

TBD

8. Normative References

[I-D.dskc-bess-bgp-car] Rao, D., Agrawal, S., Filsfils, C., Talaulikar, K., Steinberg, D., Jalil, L., Su, Y., Decraene, B., Guichard, J., Patel, K., and H. Wang, "BGP Color-Aware Routing (CAR)", draft-dskc-bess-bgp-car-03 (work in progress), October 2021. [I-D.ietf-lsr-flex-algo-bw-con] Hegde, S., J, W. B. A., Shetty, R., Decraene, B., Psenak, P., and T. Li, "Flexible Algorithms: Bandwidth, Delay, Metrics and Constraints", draft-ietf-lsr-flex-algo-bwcon-01 (work in progress), July 2021. [I-D.kaliraj-idr-bgp-classful-transport-planes] Vairavakkalai, K., Venkataraman, N., Rajagopalan, B., Mishra, G., Khaddam, M., Xu, X., Szarecki, R. J., and D. J. Gowda, "BGP Classful Transport Planes", draft-kalirajidr-bgp-classful-transport-planes-12 (work in progress), August 2021.

- [I-D.lp-lsr-bgp-algorithm] Yao, L. and P. Shaofu, "Advertisement of Algorithm in BGP", <u>draft-lp-lsr-bgp-algorithm-00</u> (work in progress), March 2021.
- [I-D.ssangli-idr-bgp-generic-metric-aigp] Sangli, S., Hegde, S., Das, R., and B. Decraene, "Generic Metric for the AIGP attribute", <u>draft-ssangli-idr-bgp-</u> <u>generic-metric-aigp-01</u> (work in progress), July 2021.
- [I-D.zhou-idr-inter-domain-lcu] Chen, R., Dai, C., and S. Peng, "Inter-domain Network Slicing via BGP-LU", <u>draft-zhou-idr-inter-domain-lcu-02</u> (work in progress), January 2021.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <u>BCP 14</u>, <u>RFC 2119</u>, DOI 10.17487/RFC2119, March 1997, <<u>https://www.rfc-editor.org/info/rfc2119</u>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", <u>RFC 5305</u>, DOI 10.17487/RFC5305, October 2008, <<u>https://www.rfc-editor.org/info/rfc5305</u>>.
- [RFC7311] Mohapatra, P., Fernando, R., Rosen, E., and J. Uttaro, "The Accumulated IGP Metric Attribute for BGP", <u>RFC 7311</u>, DOI 10.17487/RFC7311, August 2014, <<u>https://www.rfc-editor.org/info/rfc7311</u>>.
- [RFC7911] Walton, D., Retana, A., Chen, E., and J. Scudder, "Advertisement of Multiple Paths in BGP", <u>RFC 7911</u>, DOI 10.17487/RFC7911, July 2016, <<u>https://www.rfc-editor.org/info/rfc7911</u>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in <u>RFC</u> 2119 Key Words", <u>BCP 14</u>, <u>RFC 8174</u>, DOI 10.17487/RFC8174, May 2017, <<u>https://www.rfc-editor.org/info/rfc8174</u>>.
- [RFC8570] Ginsberg, L., Ed., Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", <u>RFC 8570</u>, DOI 10.17487/RFC8570, March 2019, <<u>https://www.rfc-editor.org/info/rfc8570</u>>.

Authors' Addresses

Shaofu Peng ZTE Corporation China

Internet-Draft

Email: peng.shaofu@zte.com.cn

Bin Tan ZTE Corporation China

Email: tan.bin@zte.com.cn