Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 12, 2019

L. Ginsberg, Ed.
Cisco Systems, Inc.
S. Previdi
Q. Wu
Huawei
J. Tantsura
Apstra, Inc.
C. Filsfils
Cisco Systems, Inc.

October 9, 2018

BGP-LS Advertisement of IGP Traffic Engineering Performance Metric Extensions
draft-ietf-idr-te-pm-bgp-13

Abstract

This document defines new BGP-LS TLVs in order to carry the IGP Traffic Engineering Extensions defined in IS-IS and OSPF protocols.

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.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

BGP-LS ([RFC7752]) defines NLRI and attributes in order to carry link-state information. New BGP-LS Link-Attribute TLVs are required in order to carry the Traffic Engineering Metric Extensions defined in [RFC7810] and [RFC7471].

2. Link Attribute TLVs for TE Metric Extensions

The following new Link Attribute TLVs are defined:

TLV Name
Unidirectional Link Delay
Min/Max Unidirectional Link Delay
Unidirectional Delay Variation
Unidirectional Link Loss
Unidirectional Residual Bandwidth
Unidirectional Available Bandwidth
Unidirectional Bandwidth Utilization

3. TLV Details

3.1. Unidirectional Link Delay TLV

This TLV advertises the average link delay between two directly connected IGP link-state neighbors. The semantic of the TLV is described in [RFC7810] and [RFC7471].

```
 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              |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|A|  RESERVED   |                   Delay                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Figure 1

Type: 1114

Length: 4.

3.2. Min/Max Unidirectional Link Delay TLV

This sub-TLV advertises the minimum and maximum delay values between two directly connected IGP link-state neighbors. The semantic of the TLV is described in [RFC7810] and [RFC7471].

```
     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            |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |A| RESERVED    |                   Min Delay                   |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |   RESERVED    |                   Max Delay                   |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Figure 2

Type: 1115

Length: 8.

3.3. Unidirectional Delay Variation TLV

This sub-TLV advertises the average link delay variation between two directly connected IGP link-state neighbors. The semantic of the TLV is described in [RFC7810] and [RFC7471].

```
     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     |               Delay Variation                 |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Figure 3

Type: 1116

Length: 4.

3.4. Unidirectional Link Loss TLV

This sub-TLV advertises the loss (as a packet percentage) between two directly connected IGP link-state neighbors. The semantic of the TLV is described in [RFC7810] and [RFC7471].

```
     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              |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |A|  RESERVED   |                 Link Loss                     |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Type:1117

Length: 4.

3.5. Unidirectional Residual Bandwidth TLV

This sub-TLV advertises the residual bandwidth between two directly connected IGP link-state neighbors. The semantic of the TLV is described in [RFC7810] and [RFC7471].

```
     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              |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                          Residual Bandwidth                   |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Type: 1118

Length: 4.

3.6. Unidirectional Available Bandwidth TLV

This sub-TLV advertises the available bandwidth between two directly connected IGP link-state neighbors. The semantic of the TLV is described in [RFC7810] and [RFC7471].

```
     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              |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                      Available Bandwidth                      |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Figure 4

Type: 1119

Length: 4.

3.7. Unidirectional Utilized Bandwidth TLV

This sub-TLV advertises the bandwidth utilization between two directly connected IGP link-state neighbors. The semantic of the TLV is described in [RFC7810] and [RFC7471].

```
     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              |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                     Utilized Bandwidth                        |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

where:

Figure 5

Type: 1120

Length: 4.

4. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the 'Security Considerations' section of [RFC4271] for a discussion of BGP security. Also refer to [RFC4272] and [RFC6952] for analysis of security issues for BGP.

The TLVs introduced in this document are used to propagate IGP defined information ([RFC7810] and [RFC7471].) These TLVs represent the state and resources availability of the IGP link. The IGP

instances originating these TLVs are assumed to have all the required security and authentication mechanism (as described in [RFC7810] and [RFC7471]) in order to prevent any security issue when propagating the TLVs into BGP-LS.

5. IANA Considerations

This document requests assigning code-points from the registry "BGP-LS Node Descriptor, Link Descriptor, Prefix Descriptor, and Attribute TLVs" for the new Link Attribute TLVs defined in the table below:

TLV code-point	Value
1114	Unidirectional Link Delay
1115	Min/Max Unidirectional Link Delay
1116	Unidirectional Delay Variation
1117	Unidirectional Link Loss
1118	Unidirectional Residual Bandwidth
1119	Unidirectional Available Bandwidth
1120	Unidirectional Bandwidth Utilization

6. Contributors

The following people have substantially contributed to this document and should be considered co-authors:

Saikat Ray
Individual
Email: raysaikat@gmail.com

Hannes Gredler
RtBrick Inc.
Email: hannes@rtbrick.com

7. Acknowledgements

The authors wish to acknowledge comments from Ketan Talaulikar.

8. References

8.1. 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>.

[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>.

[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>.

[RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <https://www.rfc-editor.org/info/rfc7752>.

[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>.

[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>.

8.2. Informative References

[RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <https://www.rfc-editor.org/info/rfc4272>.

[RFC6952] Jethanandani, M., Patel, K., and L. Zheng, "Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6952, DOI 10.17487/RFC6952, May 2013, <https://www.rfc-editor.org/info/rfc6952>.

Authors' Addresses

Les Ginsberg (editor)
Cisco Systems, Inc.
US

Email: ginsberg@cisco.com

Stefano Previdi
Huawei
IT

Email: stefano@previdi.net

Qin Wu
Huawei
101 Software Avenue, Yuhua District
Nanjing, Jiangsu 210012
China

Email: bill.wu@huawei.com

Jeff Tantsura
Apstra, Inc.
US

Email: jefftant.ietf@gmail.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels
BE

Email: cfilsfil@cisco.com