Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 2, 2016

S. Previdi, Ed.
Cisco Systems, Inc.
Q. Wu
Huawei
H. Gredler
S. Ray
Individual
J. Tantsura
Ericsson
C. Filsfils
L. Ginsberg
Cisco Systems, Inc.
March 1, 2016

BGP-LS Traffic Engineering (TE) Metric Extensions
draft-previdi-idr-bgpls-te-metric-extensions-01

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", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying RFC-2119 significance.

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 http://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 2, 2016.

Table of Contents

1. Introduction

BGP-LS ([I-D.ietf-idr-ls-distribution]) 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 [I-D.ietf-isis-te-metric-extensions] and [RFC7471].

2. Link Attribute TLVs for TE Metric Extensions

The following new Link Attribute TLVs are defined:

TLV Type	Value
1104 (Suggested)	Unidirectional Link Delay
1105 (Suggested)	Min/Max Unidirectional Link Delay
1106 (Suggested)	Unidirectional Delay Variation
1107 (Suggested)	Unidirectional Packet Loss
1108 (Suggested)	Unidirectional Residual Bandwidth
1109 (Suggested)	Unidirectional Available Bandwidth
1110 (Suggested)	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 [I-D.ietf-isis-te-metric-extensions] and [RFC7471].

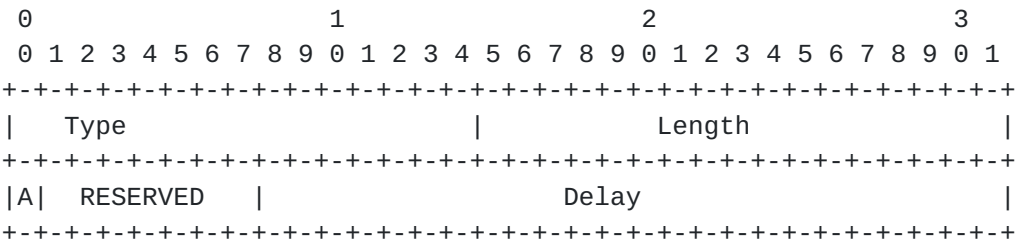

where:

Figure 1

Type: TBA (suggested value: 1104).

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 [I-D.ietf-isis-te-metric-extensions] and [RFC7471].

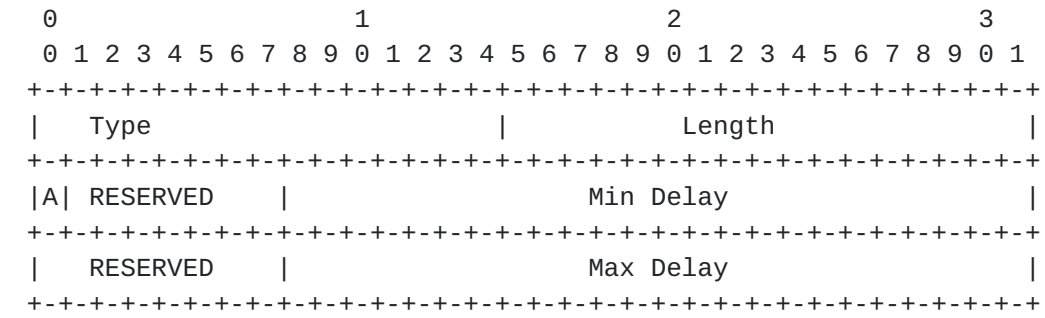

where:

Figure 2

Type: TBA (suggested value: 1105).

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 [I-D.ietf-isis-te-metric-extensions] and [RFC7471].

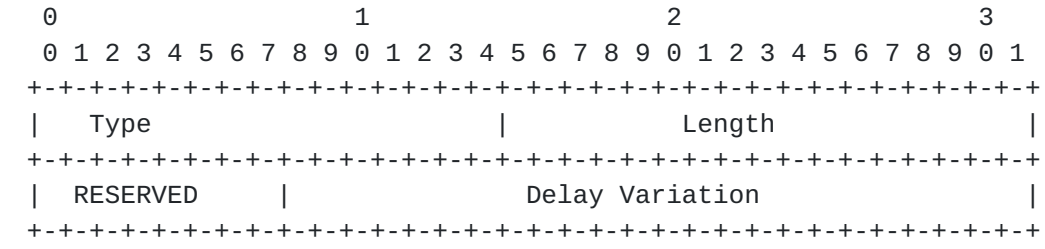

where:

Figure 3

Type: TBA (suggested value: 1106).

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 [I-D.ietf-isis-te-metric-extensions] and [RFC7471].

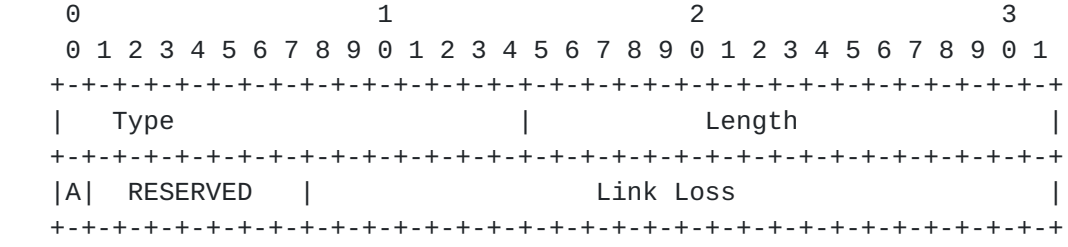

where:

Type: TBA (suggested value: 1107).

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 [I-D.ietf-isis-te-metric-extensions] and [RFC7471].

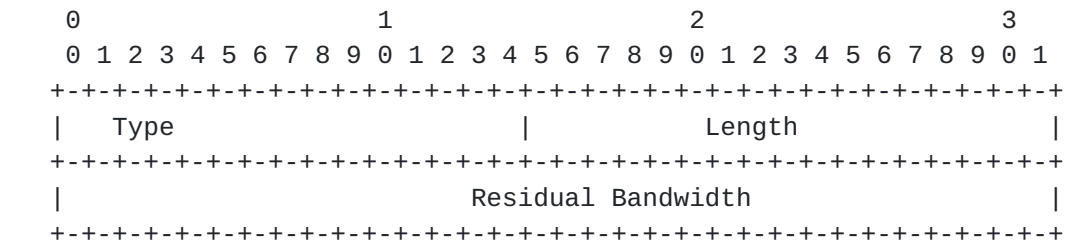

where:

Type: TBA (suggested value: 1108).

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 [I-D.ietf-isis-te-metric-extensions] and [RFC7471].

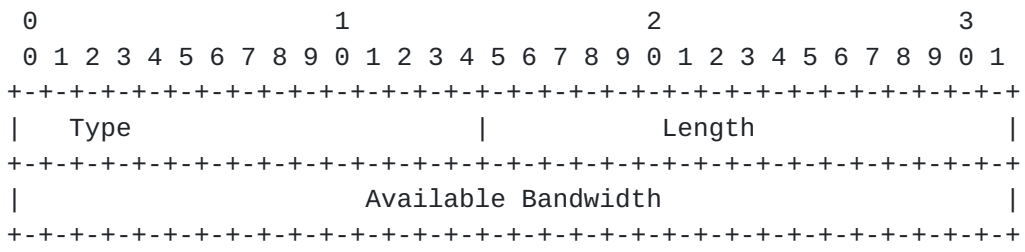

where:

Figure 4

Type: TBA (suggested value: 1109).

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 [I-D.ietf-isis-te-metric-extensions] and [RFC7471].

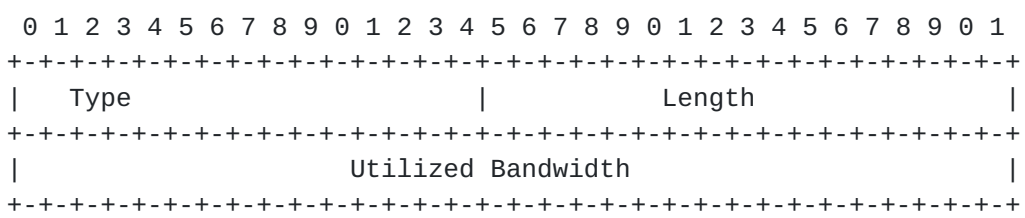

where:

Figure 5

Type: TBA (suggested value: 1110).

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 ([I-D.ietf-isis-te-metric-extensions] 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 [I-D.ietf-isis-te-metric-extensions] 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 deefined in the table here below:

TLV code-point	Value
1104 (Suggested)	Unidirectional Link Delay
1105 (Suggested)	Min/Max Unidirectional Link Delay
1106 (Suggested)	Unidirectional Delay Variation
1107 (Suggested)	Unidirectional Packet Loss
1108 (Suggested)	Unidirectional Residual Bandwidth
1109 (Suggested)	Unidirectional Available Bandwidth
1110 (Suggested)	Unidirectional Bandwidth Utilization

6. Acknowledgements

TBD

7. References

7.1. Normative References

[I-D.ietf-idr-ls-distribution]
Gredler, H., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and TE Information using BGP", draft-ietf-idr-ls-distribution-13 (work in progress), October 2015.

[I-D.ietf-isis-te-metric-extensions]
Previdi, S., Giacalone, S., Ward, D., Drake, J., and W. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", draft-ietf-isis-te-metric-extensions-11 (work in progress), February 2016.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc7471>.

7.2. Informative References

[RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <http://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, <http://www.rfc-editor.org/info/rfc6952>.

Authors' Addresses

Stefano Previdi (editor)
Cisco Systems, Inc.
Via Del Serafico 200
Rome 00191
IT

Email: sprevidi@cisco.com

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

Email: bill.wu@huawei.com

Hannes Gredler
Individual
AT

Email: hannes@gredler.at

Saikat Ray
Individual
US

Email: raysaikat@gmail.com

Jeff Tantsura
Ericsson
300 Holger Way
San Jose, CA 95134
US

Email: jeff.tantsura@ericsson.com

Clarence Filsfils
Cisco Systems, Inc.
Brussels
BE

Email: cfilsfil@cisco.com

Les Ginsberg
Cisco Systems, Inc.
US

Email: ginsberg@cisco.com