

LSVR
Internet-Draft
Intended status: Standards Track
Expires: 8 January 2026

L. Zhang
J. Dong
Huawei Technologies
7 July 2025

Applying BGP-LS Traffic Engineering Extensions to BGP-LS-SPF
draft-li-lsvr-bgp-spf-te-00

Abstract

This documents propose to introduce the BGP Link-State (BGP-LS) extensions for Traffic Engineering (TE) to the BGP-LS-SPF SAFI.

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 8 January 2026.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. Link Attribute TLVs for TE Metric Extensions	3
2.1. Administrative group(color)	4
2.2. Maximum Link Bandwidth	5
2.3. Max.reservable link bandwidth	5
2.4. Unreserved bandwidth	6
2.5. TE Default Metric	7
2.6. Link Protection Type	7
2.7. Shared Risk Link Group	8
2.8. Unidirectional Link Delay	8
2.9. Min/Max Unidirectional Link Delay	9
2.10. Unidirectional Delay Variation	9
2.11. Unidirectional Link Loss	10
2.12. Unidirectional Residual Bandwidth	10
2.13. Unidirectional Available Bandwidth	11
2.14. Unidirectional Utilized Bandwidth	11
3. Security Considerations	11
4. IANA Considerations	11
5. Acknowledgements	11
6. References	11
6.1. Normative References	11
6.2. Informative References	12
Authors' Addresses	13

1. Introduction

[I-D.ietf-lsvr-bgp-spf] extends BGP for Link-State (LS) distribution and the Shortest Path First (SPF) algorithm based calculation. BGP-LS-SPF leverages the mechanisms of both BGP protocol [RFC4271] and BGP-LS protocol extensions [RFC9552], with the extensions to BGP-LS attribute and new NLRI selection rules.

BGP-LS-SPF may be applied to network scenarios beyond data center(Such as WAN). In some network scenarios, traffic engineering is necessary to improve the resource utilization rate and load balancing. This document proposes to introduce the BGP Link-State (BGP-LS) extensions for Traffic Engineering(TE) to the BGP-LS-SPF SAFI, and discusses which TE extensions can be applied to BGP-LS-SPF SAFI.

1.1. 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.

2. Link Attribute TLVs for TE Metric Extensions

Section 5.3.2 of [RFC9552] defines the Link Attributes TLV for BGP-LS, which includes the basic TE attributes TLV. Furthermore, [RFC8571] extends the link attribute TLVs for TE, and newly defines 7 TE link attribute TLVs. The TE link attribute TLVs that can be applied to BGP-LS-SPF are shown as follows:

Type	Description	Reference
1088	Administrative group(color)	RFC 9552
1089	Maximum link bandwidth	RFC 9552
1090	Max.reservable link bandwidth	RFC 9552
1091	Unreserved bandwidth	RFC 9552
1092	TE Default Metric	RFC 9552
1093	Link Protection Type	RFC 9552
1096	Shared Risk Link Group	RFC 9552
1114	Unidirectional Link Delay	RFC 9552
1115	Min/Max Unidirectional Link Delay	RFC 9552
1116	Unidirectional Delay Variation	RFC 9552
1117	Unidirectional Link Loss	RFC 9552
1118	Unidirectional Residual Bandwidth	RFC 9552
1119	Unidirectional Available Bandwidth	RFC 9552
1120	Unidirectional Utilized Bandwidth	RFC 9552

Table 1: BGP-LS link attribute TLVs for TE metric extensions

2.1. Administrative group(color)

The administrative group sub-TLV contains a 4-octet bit mask assigned by the network administrator. The format of administrative group TLV of BGP-LS-SPF is consistent with that in BGP-LS. The format of it is shown as follow:




Figure 1: Format of administrative group TLV

where:

Bit mask: 32-bit length, each set bit corresponds to one administrative group assigned to the interface. The least significant bit is referred to as 'group 0', and the most significant bit is referred to as 'group 31'.

2.2. Maximum Link Bandwidth

The maximum link bandwidth TLV describes the maximum bandwidth that can be used on this link in this direction. This is useful for traffic engineering. The format of maximum link bandwidth TLV of BGP-LS-SPF is consistent with that in BGP-LS. The format of it is shown as follow:

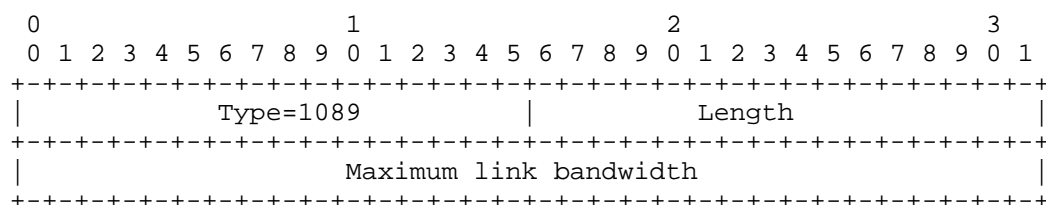


Figure 2: Format of maximum link bandwidth TLV

where:

Maximum link bandwidth: 32-bit length, it is encoded in 32 bits in IEEE floating point format. The units are bytes per second.

2.3. Max.reservable link bandwidth

The max.reservable link bandwidth TLV describes the maximum amount of bandwidth that can be reserved in this direction on this link. For oversubscription purposes, this can be greater than the bandwidth of the link. The format of max.reservable link bandwidth TLV of BGP-LS-SPF is consistent with that in BGP-LS. The format of it is shown as follow:

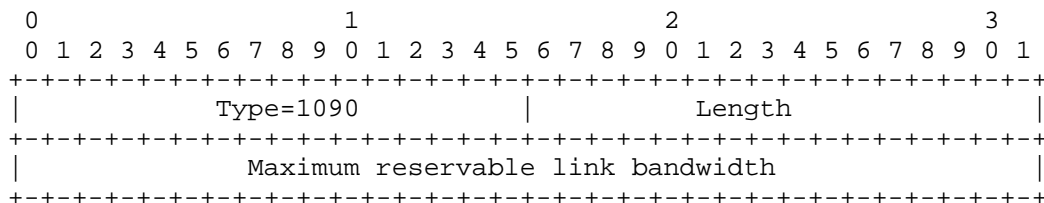


Figure 3: Format of Max.reservable link bandwidth TLV

where:

Maximum reservable link bandwidth: 32-bit length, it is encoded in 32 bits in IEEE floating point format. The units are bytes per second.

2.4. Unreserved bandwidth

The unreserved bandwidth TLV describes the amount of bandwidth reservable in this direction on this link. For oversubscription purposes, this can be greater than the bandwidth of the link. The format of unreserved bandwidth TLV of BGP-LS-SPF is consistent with that in BGP-LS. The format of it is shown as follow:




Figure 4: Format of unreserved bandwidth TLV

where:

Unreserved bandwidth(0-8): 32-bit length for each, each is encoded in 32 bits in IEEE floating point format. The units are bytes per second. The values correspond to the bandwidth that can be reserved with a setup priority of 0 through 7, arranged in increasing order with priority 0 occurring at the start of the TLV, and priority 7 at the end of the TLV.

For stability reasons, rapid changes in the values in this TLV SHOULD NOT cause rapid generation of BGP update messages.

2.5. TE Default Metric

The TE Default Metric TLV describes the Traffic Engineering metric for this link. This metric is administratively assigned and can be used to present a differently weighted topology to traffic engineering SPF calculations. The format of TE Default Metric TLV of BGP-LS-SPF is consistent with that in BGP-LS. The format of it is shown as follow:

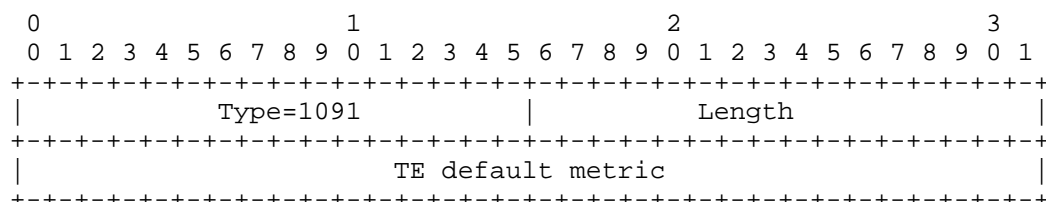


Figure 5: Format of TE Default Metric TLV

where:

TE default metric: 32-bit length metric value.

2.6. Link Protection Type

The link protection type TLV describes the protection capabilities of the link.

The format of Link Protection Type TLV of BGP-LS-SPF is consistent with that in BGP-LS. The format of it is shown as follow:

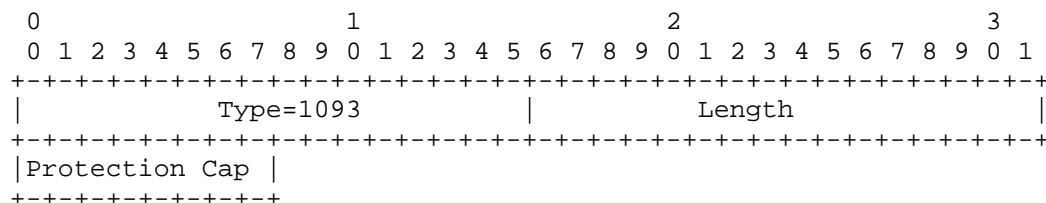


Figure 6: Format of Link Protection Type TLV

where:

Protection Cap: 8-bit length, indicates the protection capabilities of the link, for the detailed description, see Section 1.2 of [RFC5307].

2.7. Shared Risk Link Group

The Shared Risk Link Group (SRLG) TLV carries the Shared Risk Link Group information. The format of Shared Risk Link Group TLV of BGP-LS-SPF is consistent with that in BGP-LS. The format of it is shown as follow:

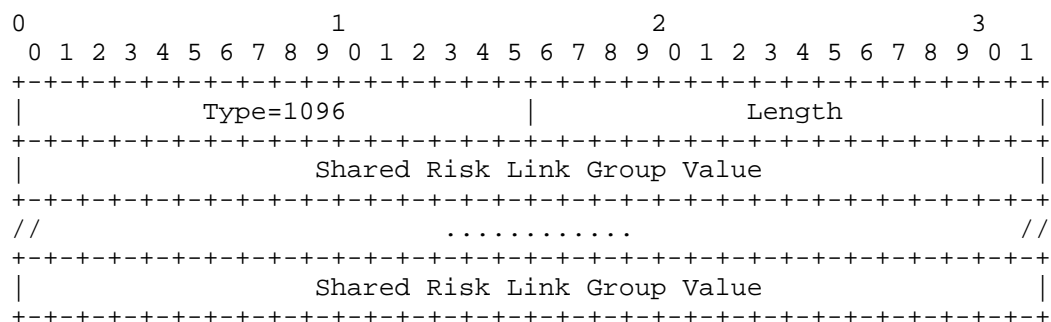


Figure 7: Format of Shared Risk Link Group TLV

where:

Shared Risk Link Group Value: variable, consisting of a (variable) list of SRLG values, where each element in the list has 4 octetslength.

2.8. Unidirectional Link Delay

This TLV describes the average link delay between two directly connected BGP-LS-SPF neighbors. The format of Unidirectional Link Delay TLV of BGP-LS-SPF is consistent with that in BGP-LS. The format of it is shown as follow:

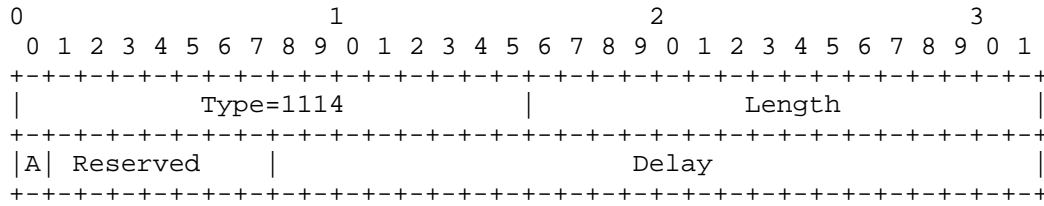


Figure 8: Format of Unidirectional Link Delay TLV

where:

A bit: This field represents the Anomalous (A) bit. For detail, see Section 4.1 of [RFC8750].

Delay: 24-bit field indicates the average link delay over a configurable interval in microseconds, encoded as an integer value.

2.9. Min/Max Unidirectional Link Delay

The Min/Max Unidirectional Link Delay TLV indicates the minimum and maximum delay values between two directly connected BGP-LS-SPF neighbors. The semantics and values of the fields in the TLV are the same as that described in [RFC8570] and [RFC7471].




Figure 9: Format of Min/Max Unidirectional Link Delay TLV

2.10. Unidirectional Delay Variation

The Unidirectional Delay Variation describes the average link delay variation between two directly connected BGP-LS-SPF neighbors. The semantics and values of the fields in the TLV are the same as that described in [RFC8570] and [RFC7471].




Figure 10: Format of Unidirectional Delay Variation TLV

2.11. Unidirectional Link Loss

This TLV describes the loss (as a packet percentage) between two directly connected BGP-LS-SPF neighbors. The semantics and values of the fields in the TLV are the same as that described in [RFC8570] and [RFC7471].

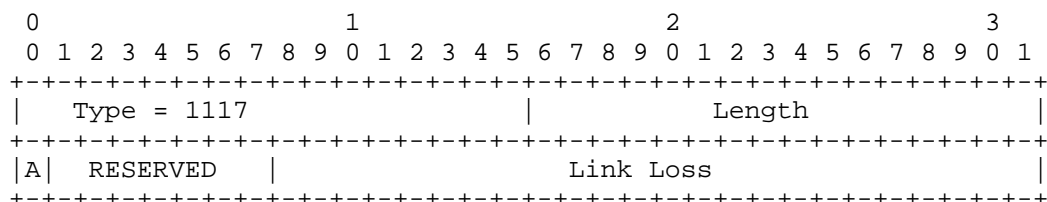


Figure 11: Format of Unidirectional Link Loss TLV

2.12. Unidirectional Residual Bandwidth

This TLV advertises the residual bandwidth between two directly connected BGP-LS-SPF neighbors. The semantics and values of the fields in the TLV are the same as that described in [RFC8570] and [RFC7471].

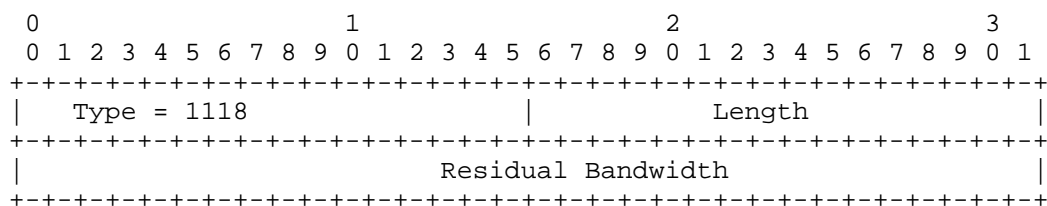


Figure 12: Format of Unidirectional Residual Bandwidth TLV

2.13. Unidirectional Available Bandwidth

This TLV advertises the available bandwidth between two directly connected BGP-LS-SPF neighbors. The semantics and values of the fields in the TLV are the same as that described in [RFC8570] and [RFC7471].

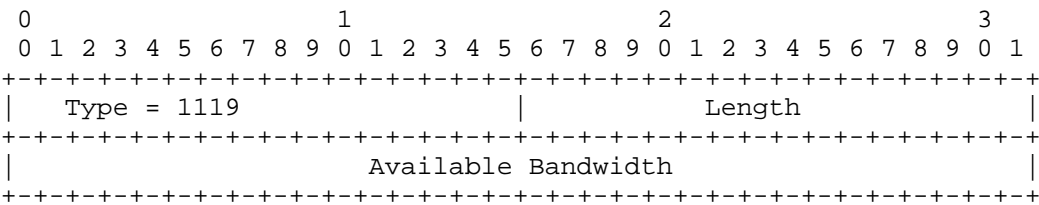


Figure 13: Format of Unidirectional Residual Bandwidth TLV

2.14. Unidirectional Utilized Bandwidth

This TLV advertises the bandwidth utilization between two directly connected BGP-LS-SPF neighbors. The semantics and values of the fields in the TLV are the same as that described in [RFC8570] and [RFC7471].

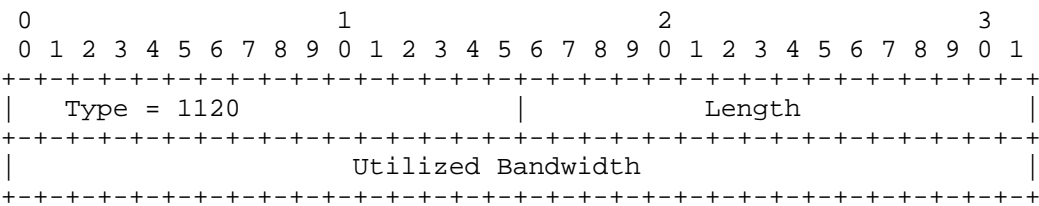


Figure 14: Format of Unidirectional Utilized Bandwidth TLV

3. Security Considerations

This document introduces no additional security vulnerabilities in addition to the ones as described in [RFC9552] and [RFC8571].

4. IANA Considerations

This document has no IANA actions.

5. Acknowledgements

6. References

6.1. Normative References

[I-D.ietf-lsvr-bgp-spf]

Patel, K., Lindem, A., Zandi, S., and W. Henderickx, "BGP Link-State Shortest Path First (SPF) Routing", Work in Progress, Internet-Draft, draft-ietf-lsvr-bgp-spf-51, 23 January 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-lsvr-bgp-spf-51>>.

[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>>.

[RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.

[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>>.

6.2. Informative References

[RFC8571] Ginsberg, L., Ed., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions", RFC 8571, DOI 10.17487/RFC8571, March 2019, <<https://www.rfc-editor.org/info/rfc8571>>.

[RFC5307] Kompella, K., Ed. and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, DOI 10.17487/RFC5307, October 2008, <<https://www.rfc-editor.org/info/rfc5307>>.

[RFC8750] Migault, D., Guggemos, T., and Y. Nir, "Implicit Initialization Vector (IV) for Counter-Based Ciphers in Encapsulating Security Payload (ESP)", RFC 8750, DOI 10.17487/RFC8750, March 2020, <<https://www.rfc-editor.org/info/rfc8750>>.

- [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>>.
- [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>>.

Authors' Addresses

Li Zhang
Huawei Technologies
No. 156 Beiqing Road
Beijing
China
Email: zhangli344@huawei.com

Jie Dong
Huawei Technologies
No. 156 Beiqing Road
Beijing
China
Email: jie.dong@huawei.com