Network Working Group
Internet Draft
Intended status: Standards Track
Expires: August 13, 2026

C. Lin
New H3C Technologies
Y. Liu
China Mobile
Y. Zhu
China Telecom
R. Chen
ZTE Corporation

February 09, 2026

BGP-LS Advertisement of SR Policy Performance Metric
draft-lin-idr-bgpls-te-policy-pm-08

Abstract

This document describes a way to advertise the performance metrics for Traffic Engineering (TE) Policy using BGP Link State (BGP-LS).

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

This Internet-Draft will expire on August 13, 2026.

Table of Contents

1. Introduction

BGP Link State (BGP-LS) can be used to distribute link-state and traffic engineering (TE) information to external components [RFC9552]. [RFC9857] describes the mechanism for BGP-LS to distribute the information of TE policies.

In some network scenarios, the controller needs to obtain the performance information of TE Policies, which can be used in service placement to meet better customer requirements and utilize network resources more efficiently.

[I-D.ietf-spring-stamp-srpm-mpls] and [I-D.ietf-spring-stamp-srpm-srv6] describe the procedures for Performance Measurement in SR

networks, using STAMP as defined in [RFC8762]. The described procedure is used for SR paths [RFC8402] (including SR Policies [RFC9256]).

This document describes a way to advertise the performance metrics for Traffic Engineering (TE) Policy using BGP Link State (BGP-LS).

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. Advertisement of SR Policy Performance Metric

[RFC8571] defines several Link Attribute TLVs for BGP-LS to carry the IGP Traffic Engineering Performance Metric Extensions:

TLV Code Point	Value
1114	Unidirectional Link Delay
1115	Min/Max Unidirectional Link Delay
1116	Unidirectional Delay Variation
1117	Unidirectional Link Loss

The above TLVs can be reused to report performance metrics for TE Policies. They enable reporting for either a specific SID-List or a particular SR candidate path. Therefore, these TLVs may serve as sub-TLVs of the SR Segment List TLV, or alternatively provide performance metrics for the SR Policy at the candidate path level.

When used to describe the performance metric of the SR Policy NLRI, they are carried in the optional non-transitive BGP Path Attribute "BGP-LS Attribute" defined in [RFC9552]. The semantics of the above TLVs comply with [RFC8571], except for that they are extended to describe TE Policies besides IGP links.

The performance metric of SR Policy may be measured at the headend, for example, by using STAMP for SR Policy [I-D.ietf-spring-stamp-srpm-mpls] and [I-D.ietf-spring-stamp-srpm-srv6]. But the measurement methods are out of the scope of this document.

The existing performance metrics above are all unidirectional. However, there are also requirements to advertise round-trip

performance metrics for TE Policies. The BGP-LS extensions for round-trip TE performance metrics are defined in the following section.

3. Extensions for Round-trip TE Performance Metric

3.1. Round-trip Delay TLV

This TLV advertises the average round-trip delay for SR Policy. It can report round-trip delay for either a specific SID-List or a particular SR candidate path. As such, this TLV may function either as a sub-TLV of the SR Segment List TLV or provide round-trip delay for the SR Policy at the candidate path level.

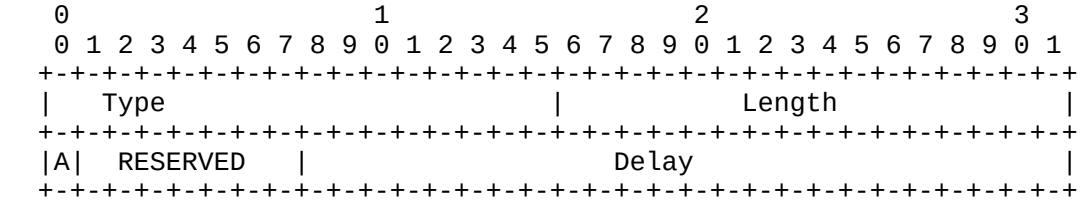

where:

- Type: TBD
- Length: 4
- Reserved: Reserved for future use. MUST be set to 0 when sent and MUST be ignored when received.
- A: Anomalous (A) Bit. Same with the A Bit in Unidirectional Link Delay TLV [RFC8571].
- Delay: Similar with the Delay filed in Unidirectional Link Delay TLV [RFC8571], except for that the delay is round-trip.

3.2. Min/Max Round-trip Delay TLV

This TLV advertises the minimum and maximum round-trip delay for SR Policy. It can report minimum and maximum round-trip delay for either a specific SID-List or a particular SR candidate path. As such, this TLV may function either as a sub-TLV of the SR Segment List TLV or provide minimum and maximum round-trip delay for the SR Policy at the candidate path level.

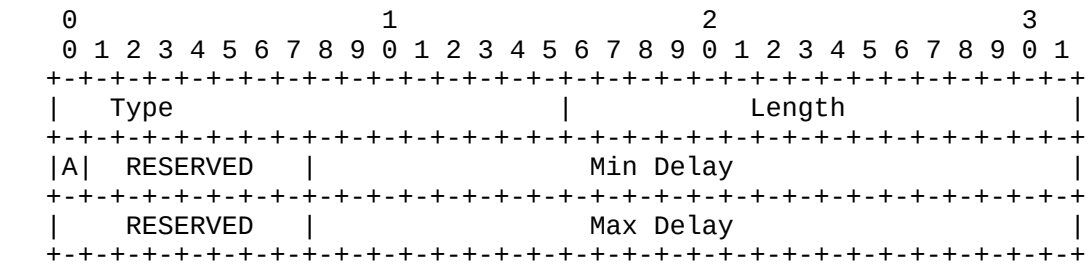

where:

- Type: TBD
- Length: 4
- Reserved: Reserved for future use. MUST be set to 0 when sent and MUST be ignored when received.
- A: Anomalous (A) Bit. Same with the A Bit in Min/Max Unidirectional Link Delay TLV [RFC8571].
- Min Delay: Similar with the Min Delay filed in Min/Max Unidirectional Link Delay TLV [RFC8571], except for that the delay is round-trip.
- Max Delay: Similar with the Max Delay filed in Min/Max Unidirectional Link Delay TLV [RFC8571], except for that the delay is round-trip.

3.3. Round-trip Delay Variation TLV

This TLV advertises the average round-trip delay variation for SR Policy. It can report average round-trip delay variation for either a specific SID-List or a particular SR candidate path. As such, this TLV may function either as a sub-TLV of the SR Segment List TLV or provide average round-trip delay variation for the SR Policy at the candidate path level.

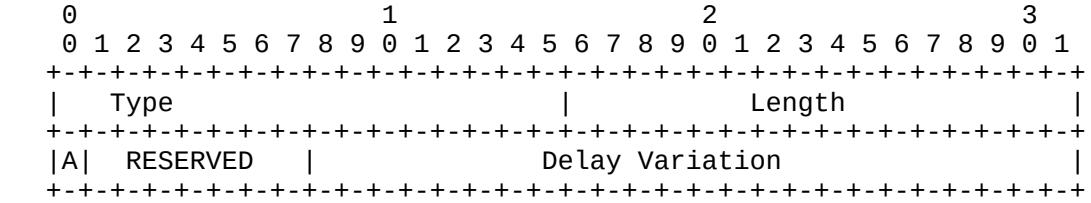

where:

- Type: TBD
- Length: 4
- Reserved: Reserved for future use. MUST be set to 0 when sent and MUST be ignored when received.
- A: Anomalous (A) Bit. Same with the A Bit in Unidirectional Delay Variation TLV [RFC8571].
- Delay Variation: Similar with the Delay Variation filed in Unidirectional Delay Variation TLV [RFC8571], except for that the delay variation is round-trip.

3.4. Round-trip Loss TLV

This TLV advertises the round-trip loss for SR Policy. It can report round-trip loss for either a specific SID-List or a particular SR candidate path. As such, this TLV may function either as a sub-TLV of the SR Segment List TLV or provide round-trip loss for the SR Policy at the candidate path level.

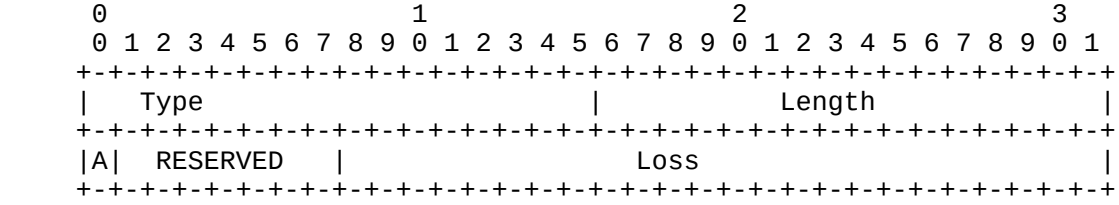

where:

- Type: TBD
- Length: 4
- Reserved: Reserved for future use. MUST be set to 0 when sent and MUST be ignored when received.
- A: Anomalous (A) Bit. Same with the A Bit in Unidirectional Link Loss TLV [RFC8571].
- Loss: Similar with the Link Loss filed in Unidirectional Link Loss TLV [RFC8571], except for that the loss is round-trip.

4. Security Considerations

The procedures and protocol extensions defined in this document do not affect the BGP security model.

The sub-TLVs introduced in this document allow an operator to advertise state information of SR Policy (e.g., bandwidth, delay), which may be sensitive and dynamic in nature.

In very large networks, instability could occur if measurement intervals are configured at a frequency that overwhelms the processing capacity for announcement intervals. Therefore, care must be taken when configuring these values. Implementations SHOULD NOT allow the inter-update timer to be set lower than the measurement interval. Additionally, implementations SHOULD enforce configurable constraints to mitigate the risk of instability.

For a discussion of BGP security, refer to the "Security Considerations" section of [RFC4271]. Analyses of BGP security issues are also provided in [RFC4272] and [RFC6952]. Security considerations related to the acquisition and distribution of BGP-LS information are discussed in [RFC9552], [RFC9830], and [RFC9857].

The mechanism proposed in this document is subject to the same vulnerabilities as any other protocol that relies on BGP-LS.

5. Management Considerations

An implementation SHOULD allow the operator to specify neighbors to which Link-State NLRIs will be advertised and from which Link-State NLRIs will be accepted.

An implementation SHOULD allow the operator to control the content of advertisements, such as whether or not to advertise latency, packet loss rate, bidirectional latency, and bidirectional packet loss rate.

An implementation SHOULD allow the operator to control advertisement thresholds to avoid frequent announcements.

6. IANA Considerations

This document defines the following TLVs for BGP-LS.

TLV Code Point	Value
TBD	Round-trip Delay
TBD	Min/Max Round-trip Delay
TBD	Round-trip Variation
TBD	Round-trip Loss

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[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>.

[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>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017

[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>.

[RFC9830] Previdi, S., Filsfils, C., Talaulikar, K., Mattes, P., and Jain, D., "Advertising Segment Routing Policies in BGP", RFC 9857, DOI 10.17487/RFC9830, September 2025,

[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>.

[RFC9857] Previdi, S., Talaulikar, K., Dong, J., Gredler, H., and Tantsura, J., "Advertisement of Segment Routing Policies using BGP Link-State", RFC 9857, DOI 10.17487/RFC9857, October 2025, <https://www.rfc-editor.org/info/rfc9857>.

7.2. Informative References

[I-D.ietf-spring-stamp-srpm-mpls] Gandhi, R., Filsfils, C., Voyer, D., Chen, M., and R. F.Foote, "Performance Measurement Using Simple Two-Way Active Measurement Protocol (STAMP) for Segment Routing over the MPLS Data Plane", Work in Progress, Internet-Draft, draft-ietf-spring-stamp-srpm-mpls-00, 02 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-stamp-srpm-mpls-00>.

[I-D.ietf-spring-stamp-srpm-srv6] Gandhi, R., Filsfils, C., Voyer, D., Chen, M., and R. F.Foote, "Performance Measurement Using Simple Two-Way Active Measurement Protocol (STAMP) for Segment Routing over the SRv6 Data Plane ", Work in Progress, Internet-Draft, draft-ietf-spring-stamp-srpm-srv6-00, 02 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-stamp-srpm-srv6-00>.

Authors' Addresses

Changwang Lin
New H3C Technologies

Email: linchangwang.04414@h3c.com

Yisong Liu
China Mobile

Email: liuyisong@chinamobile.com

Yongqing Zhu
China Telecom
Guangzhou
Email: zhuyq8@chinatelecom.cn

Ran Chen
ZTE Corporation
China
Email: chen.ran@zte.com.cn