

LSR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 23 January 2026

R. Gandhi, Ed.
Cisco Systems, Inc.
22 July 2025

IS-IS, OSPF and BGP-LS Traffic Engineering and Flexible Algorithm
Extensions for Utilizing Bit Error Rate Metrics
draft-gandhi-lsr-ber-00

Abstract

This document describes extensions to IS-IS, OSPF, and BGP-LS Traffic Engineering to distribute the Bit Error Rate (BER) and Packet Error Rate (PER) metrics for the links that can be used for flexible algorithm and path selection purpose.

Note that this document only covers the mechanisms with which network-performance information is distributed. The mechanisms for measuring network performance or acting on that information, once distributed, are outside the scope of this document.

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 23 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
2. Conventions Used in This Document	3
2.1. Requirements Language	3
2.2. Abbreviations	3
3. Overview	3
4. IS-IS Extensions	4
5. OSPF Extensions	5
6. BGP-LS Extensions	6
7. Flexible Algorithm Type Extensions	6
8. Security Considerations	7
9. IANA Considerations	7
10. References	8
10.1. Normative References	8
10.2. Informative References	8
Acknowledgments	9
Author's Address	9

1. Introduction

Networks may experience transmission bit errors due to various factors, such as poor fiber quality. The bit error can be a single bit error or a burst of bit errors at a time. Bit errors include layer-2 bit errors (e.g., causing CRC errors) or layer-3 and layer-4 bit errors (e.g., causing checksum failures). It is feasible to measure the Bit Error Rate (BER) and Packet Error Rate (PER) of the links using measurement packets. However, the measured BER and PER metrics are currently not used in flexible-algorithm and traffic engineering as they are not advertised in IS-IS, OSPF or BGP-LS.

[RFC8570] for IS-IS, [RFC7471] for OSPF, and [RFC8571] for BGP-LS, define TE metrics extensions for distributing latency, loss, and bandwidth metrics. These documents, however, do not define extensions for distributing BER and PER metrics.

In this document, IS-IS, OSPF, and BGP-LS extensions are defined to advertise the BER and PER metrics for the links in the network that can be used for flexible algorithm and TE. Also, IS-IS, OSPF and BGP-LS extensions are defined to advertise anomaly state of the links that can be used to avoid the links with anomaly state.

2. Conventions Used in This Document

2.1. 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.

2.2. Abbreviations

BER: Bit Error Rate

EMA: Exponential Moving Average

MTU: Maximum Transmission Unit

PER: Rate of Packets with Bit Errors

TLV: Type-Length-Value

3. Overview

In this document, IS-IS, OSPF, and BGP-LS extensions are defined to advertise the BER and PER metrics for the links in the network that can be used for flexible algorithm and TE. Also, IS-IS, OSPF and BGP-LS extensions are defined to advertise anomaly state of the links that can be used to avoid the links with anomaly state.

Various BER metrics are distributed including:

- (1) Average BER
- (2) Maximum BER
- (3) Minimum BER
- (4) Exponential Moving Average of BER
- (5) Variance of BER (difference of minimum and average, for example)
- (6) BER anomaly state

Various PER metrics are distributed including:

- (1) Average PER

- (2) Maximum PER
- (3) Minimum PER
- (4) Exponential Moving Average (EMA) of PER
- (5) Variance of PER metrics (difference of minimum and average, for example)
- (6) PER anomaly state

Editor's note: The PER metric can be added in the existing packet loss metric (instead of using it as a separate metric). However, the BER metric can provide different characteristics of the network, as depending on the BER pattern (bursty or single), there can be high BER but low packet loss, and vice versa.

Editor's note: The BER and PER metrics of the links can be used as generic metrics in IS-IS, OSPF and BGP-LS.

4. IS-IS Extensions

This document extends the sub-TLV structure defined in [RFC8570]. This document registers new IS-IS TE sub-TLVs in the "Sub-TLVs for TLVs 22, 23, 141, 222, and 223" registry. These new sub-TLVs provide ways to distribute network-performance information. The extensions in this document build on the extensions provided in IS-IS TE [RFC5305] and GMPLS [RFC4203].

The Extended IS Reachability TLV (type 22) (defined in [RFC5305]), Inter-AS Reachability TLV (also called "inter-AS reachability information TLV") (type 141) (defined in [RFC9346]), and MT-ISN TLV (type 222) (defined in [RFC5120]) have nested sub-TLVs that permit the TLVs to be readily extended. This document registers several sub-TLVs:

Type Description

TBA1 Unidirectional Link Average BER

TBA2 Unidirectional Link Average PER

Example Sub-TLV for distributing BER metric in IS-IS is shown in Figure 1.




Figure 1: IS-IS Unidirectional Average BER Metric Sub-TLV

A bit

This field represents the Anomalous (A) bit. The A bit is set when the measured value of this parameter exceeds its configured maximum threshold. The A bit is cleared when the measured value falls below its configured reuse threshold. If the A bit is cleared, the sub-TLV represents steady-state link performance.

Unit for BER and PER Metrics

This 24-bit field carries BER or PER as a percentage over a configurable interval. The basic unit is 0.000003%, where $(2^{24} - 2)$ is 50.331642%. This value is the highest percentage that can be expressed (the assumption that high-speed links with over 50% BER/PER are unusable). Measured values that are larger than the field maximum SHOULD be encoded as the maximum value.

The BER metric added is the exponential moving average (EMA) of the BER in order to suppress the frequent advertisements. The metric is only distributed when user-configured threshold is crossed for the same reason. This applies to PER metric as well.

5. OSPF Extensions

OSPF extensions for TE are defined in [RFC7471]. This document defines several additional sub-TLVs for the Link TLV:

Type Description

TBA3 Unidirectional Link Average BER

TBA4 Unidirectional Link Average PER

Example Sub-TLV for distributing BER metric in OSPF is shown in Figure 2.




Figure 2: OSPF Unidirectional Average BER Metric Sub-TLV

The A flag and Unit of BER and PER metrics are defined the same way as for IS-IS.

6. BGP-LS Extensions

The IGP/TE metrics in BGP Link-State are defined in [RFC8571]. New BGP-LS Link Attribute TLVs are defined for BER. TLV formats follow the rules defined in [RFC9552].

Value TLV Code Point

TBA5 Unidirectional Link Average BER

TBA6 Unidirectional Link Average PER

Example TLV for distributing BER metric in BGP-LS is shown in Figure 3 below.




Figure 3: BGP-LS Unidirectional Average BER Metric Sub-TLV

The A flag and Unit of BER and PER metrics are defined the same way as for IS-IS.

7. Flexible Algorithm Type Extensions

IGP Flexible algorithm is defined in RFC 9502 and RFC 9350. New IS-IS and OSPF Flexible algorithm type extensions are defined using the BER and PER metrics of the links.

Editor's note: Additional details on the procedure will be added in the future revision of this document.

8. Security Considerations

The security considerations specified in [RFC8571], [RFC8570], and [RFC7471] apply to the procedure and extensions defined in this document.

9. IANA Considerations

IANA maintains the registry for the sub-TLVs for IS-IS. IANA is requested to allocate the following sub-TLVs in the "Sub-TLVs for TLVs 22, 23, 141, 222, and 223" registry:

Value	Description	Reference
TBA1	Unidirectional Link Average BER Metric	This document
TBA2	Unidirectional Link Average PER Metric	This document

Table 1: IS-IS Sub-TLV Types

IANA maintains the registry for the sub-TLVs for OSPF link TLV. IANA is requested to allocate the following sub-TLVs for the OSPF link TLV.

Value	Description	Reference
TBA3	Unidirectional Link Average BER Metric	This document
TBA4	Unidirectional Link Average PER Metric	This document

Table 2: OSPF Sub-TLV Types

IANA maintains the registry for the link attribute TLVs in BGP-LS. IANA is requested to allocate the following link attribute TLVs in BGP-LS.

Value	Description	Reference
TBA5	Unidirectional Link Average BER Metric	This document
TBA6	Unidirectional Link Average PER Metric	This document

Table 3: BGP-LS Link attribute TLV Types

Editor's note: Additional Sub-TLVs for carrying minimum, maximum, and variance of the BER and PER metrics will be defined in the future revision of this document.

10. References

10.1. 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>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.

10.2. Informative References

- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.

- [RFC4203] Kompella, K., Ed. and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, DOI 10.17487/RFC4203, October 2005, <<https://www.rfc-editor.org/info/rfc4203>>.
- [RFC9346] Chen, M., Ginsberg, L., Previdi, S., and D. Xiaodong, "IS-IS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 9346, DOI 10.17487/RFC9346, February 2023, <<https://www.rfc-editor.org/info/rfc9346>>.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.
- [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>>.

Acknowledgments

TBA.

Author's Address

Rakesh Gandhi (editor)
Cisco Systems, Inc.
Canada
Email: rgandhi@cisco.com