Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 4 September 2025

J. Dong
X. Geng
Huawei Technologies
Y. Chen
CAICT
J. Xu
G. Zhang
China Mobile
S. Yang
China Telecom

3 March 2025

IS-IS Extensions for Load Balancing Alternates
draft-dong-lsr-load-balancing-alternate-00

Abstract

IGP normally computes the shortest paths in the network based on the IGP metric, without taking the link utilization into consideration. In network scenarios where there is link degradation or failure which causes link throughput reduction, or the volume of specific traffic flows increase dramatically, congestion can happen if only the best path is used for forwarding. This document describes IS-IS extensions for the computation of alternate traffic engineering paths which can be used for traffic load balancing when the shortest path is becoming congested.

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 4 September 2025.

Table of Contents

1. Introduction

Intermediate System to Intermediate System (IS-IS) protocol normally computes the best paths in the network based on the IGP metric cost. The utilization of the network links is not taken into consideration. When there is link degradation or failure in the network which causes throughput reduction, or the volume of specific traffic flows increase dramatically, congestion can happen if only the best path is used for forwarding. This document describes IS-IS extensions for the computation of alternate paths which can be used for traffic load balancing when the shortest path is becoming congested. A new IGP application [RFC8919] called "Load Balancing Alternate (LBA)" is introduced for the computation of alternate Traffic Engineering (TE) paths, which can be used for traffic load balancing when the shortest path is becoming congested. The set of link attributes which need to be advertised for this new application are also specified.

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 BCP14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Application for Alternate Load Balancing

For the computation of alternate load-balancing paths on network nodes, a new IGP application called "Load Balancing Alternate (LBA)" is introduced in this document, and a set of application-specific link attributes can be advertised for LBA.

2.1. Application Identifier Bit Mask for LBA

[RFC8919] defines the Application-Specific Link Attributes sub-TLV, which can be used to carry application-specific link attributes in IS-IS. A new bit "B" in the Standard Application Bit Mask (SABM) field is defined for Load Balancing Alternate, as shown in the following figure:

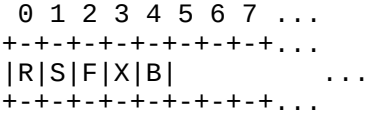

The R, S, F, X bit are defined in [RFC8919] and [RFC9350] respectively.

B-bit: When set, it indicates the application is Load-balancing Alternate.

2.2. Application-Specific Link Attributes Sub-TLV for LBA

The Application-Specific Link Attributes sub-TLV as specified in [RFC8919] is used to carry the link attributes which are used for the computation of load balancing alternates. When the B bit in the Application Identifier Bit Mask is set, the Unidirectional Utilized Bandwidth sub-TLV as defined in [RFC8570] MUST be carried as a sub-TLV under the Application-Specific Link Attributes (ASLA) sub-TLV. Other types of link attributes may also be advertised as application-specific for LBA, the details are for further study.

3. Operational Procedures

Each network node which supports the LBA application SHOULD continuously monitor the bandwidth utilization of its links, and advertise the unidirectional utilized bandwidth periodically using the Application-Specific Link Attributes sub-TLV, the advertisement interval SHOULD be configurable, and the suggested interval is 10 seconds. The utilized bandwidth information of all the network links collected from ASLA advertisement will be used as the input for the computation of load balancing alternate paths.

According to the operators' policy, the computation of Load Balancing Alternate paths can be enabled on specific network nodes. On network nodes where LBA is enabled, when the bandwidth utilization on a specific local link exceeds the pre-configured threshold (Threshold 1), the node SHOULD triggers the computation of load balancing alternate TE paths for the destinations whose next-hop of the shortest path points to this link. The collected link bandwidth utilization information SHOULD be taken into consideration in the computation. Multiple LBA paths may be computed for the same destination. The bandwidth utilization of the links on the load balancing alternate paths SHOULD meet some criteria so that themselves will not become congested when a portion of the traffic are switched from the shortest paths to the LBA paths. The LBA paths may or may not be loop-free. If an LBA path is not a loop-free alternate, some mechanism for explicit path forwarding, such as segment routing traffic engineering (SR-TE) MUST be used to ensure there is no forwarding loop.

When the bandwidth utilization on the link continues increasing and exceeds the second pre-configured threshold (Threshold 2), the network node SHOULD steer a portion of the traffic from the shortest paths to the LBA paths. If there are multiple LBA paths, the traffic load can be distributed among these paths according to the bandwidth utilization of the LBA paths. The detailed load distribution mechanism is implementation specific.

Section 4.2.3 of [RFC8919] gives the considerations for the advertisement of extended TE metrics as defined in [RFC8570], and suggests that the advertisement for these attributes is associated with all of the applications utilizing the link. While for the computation of load-balancing alternate paths, the interval of the attributes advertisement can be much shorter than the interval required by other applications. To avoid unexpected churn introduced to other applications, it is RECOMMENDED to advertise the attributes for LBA independently.

4. IANA Considerations

IANA is requested to assign a new code point from the "Link Attribute Application Identifiers" Registry.

Bit #	Description
TBA	Load balancing alternate (B)

5. Security Considerations

Procedures and protocol extensions defined in this document do not affect the security model of IS-IS. The security considerations in [RFC8570] and [RFC8919] apply to this document.

6. Acknowledgements

The authors would like to thank Mengkai Zhao and Jieyu Liang for their review and suggestions.

7. References

7.1. 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>.

[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>.

[RFC8919] Ginsberg, L., Psenak, P., Previdi, S., Henderickx, W., and J. Drake, "IS-IS Application-Specific Link Attributes", RFC 8919, DOI 10.17487/RFC8919, October 2020, <https://www.rfc-editor.org/info/rfc8919>.

7.2. Informative References

[RFC9350] Psenak, P., Ed., Hegde, S., Filsfils, C., Talaulikar, K., and A. Gulko, "IGP Flexible Algorithm", RFC 9350, DOI 10.17487/RFC9350, February 2023, <https://www.rfc-editor.org/info/rfc9350>.

Authors' Addresses

Jie Dong
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing
100095
China
Email: jie.dong@huawei.com

Xuesong Geng
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing
100095
China
Email: gengxuesong@huawei.com

Yunke Chen
CAICT
China
Email: chenyunke@caict.ac.cn

Jie Xu
China Mobile
China
Email: xujie@sc.chinamobile.com

Geng Zhang
China Mobile
China
Email: zhanggeng@chinamobile.com

Shunxing Yang
China Telecom
China
Email: yangsx@chinatelecom.cn