

LSR Working Group
Internet-Draft
Intended status: Standards Track
Expires: 1 September 2026

L. Zhang
J. Dong
C. Li
Huawei
28 February 2026

IGP Extensions for Sub-interface Relationship Information
draft-zl-lsr-igp-sub-interface-relationship-00

Abstract

This document extends ISIS and OSPF, allowing a network device to advertise the relationship between a physical interface and its sub-interfaces. This extension enables the links based on sub-interfaces to participate in the alternative paths for load balancing in SRv6 BE bandwidth polling.

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 1 September 2026.

Copyright Notice

Copyright (c) 2026 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
1.2. Terminology	3
2. ISIS extensions	3
2.1. Physical Local Link Relationship Information TLV	3
2.2. Physical Local Link Information sub-TLV	5
3. OSPF extensions	5
4. Use cases	6
5. IANA Considerations	7
6. Security Considerations	7
7. References	7
7.1. Normative References	7
7.2. Informative References	7
Acknowledgements	8
Authors' Addresses	8

1. Introduction

The SRv6 BE bandwidth pooling technology enables network devices to detect link bandwidth and bandwidth usage in real time, thereby quickly identifying faults and congestion, and automatically calculating available alternative paths to load balance traffic. This function relies on network devices advertising their interface bandwidth and bandwidth usage information to external systems. There are several RFCs defining how to advertise the maximum bandwidth and utilized bandwidth of a link by IGP.

[RFC5305] describes extensions to Intermediate System to Intermediate System (IS-IS) protocol to support Traffic Engineering (TE). It defines the Maximum Link Bandwidth sub-TLV used to describe the maximum bandwidth of a link between two directly connected IS-IS neighbors.

[RFC8750] further extends [RFC5305], and defines the Unidirectional Utilized Bandwidth Sub-TLV to describe the bandwidth utilization between two directly connected IS-IS neighbors.

[RFC5305] describes extensions to OSPF protocol to support Traffic Engineering (TE). It defines the Maximum Link Bandwidth sub-TLV in the Link TLV used to describe the maximum bandwidth of a link between two directly connected OSPF neighbors.

[RFC7471] further extends [RFC5305], and defines the Unidirectional Utilized Bandwidth Sub-TLV to describe the bandwidth utilization between two directly connected OSPF neighbors.

However, these layer 3 protocols, such as IS-IS and OSPF, often establish neighbor relationships through sub-interfaces. When two directly connected IS-IS neighbors are established based on the sub-interface, then the Maximum Link Bandwidth and Unidirectional Utilized Bandwidth of the sub-interfaces are the same as their parent physical interface, since the sub-interfaces don't have independent bandwidth and bandwidth utilization information.

The remote devices don't know the relationship between the sub-interfaces and their parent physical interface. Therefore, the remote devices don't know the maximum link bandwidth and unidirectional utilized bandwidth information is shared among all the sub-interfaces of a specific physical interface. This makes a remote link based on sub-interfaces difficult to participate in traffic load balancing as an alternative forwarding path.

This document extends ISIS and OSPF, allowing a network device to advertise the relationship between a physical interface and its sub-interfaces. This extension enables the links based on sub-interfaces to participate in the alternative paths for load balancing in SRv6 BE bandwidth polling.

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.

1.2. Terminology

2. ISIS extensions

There are two extension options to advertise the relationship between a physical interface and its sub-interfaces.

2.1. Physical Local Link Relationship Information TLV

This extension option defines a new top level TLV named Physical Local Link Relationship Information TLV. This TLV describes the relationship between a physical interface and its sub-interfaces, and the physical bandwidth and bandwidth utilization information of the physical interface. The physical bandwidth and bandwidth utilization information is shared for all the sub-interfaces.

The format of Interface Relationship TLV is shown as follows:

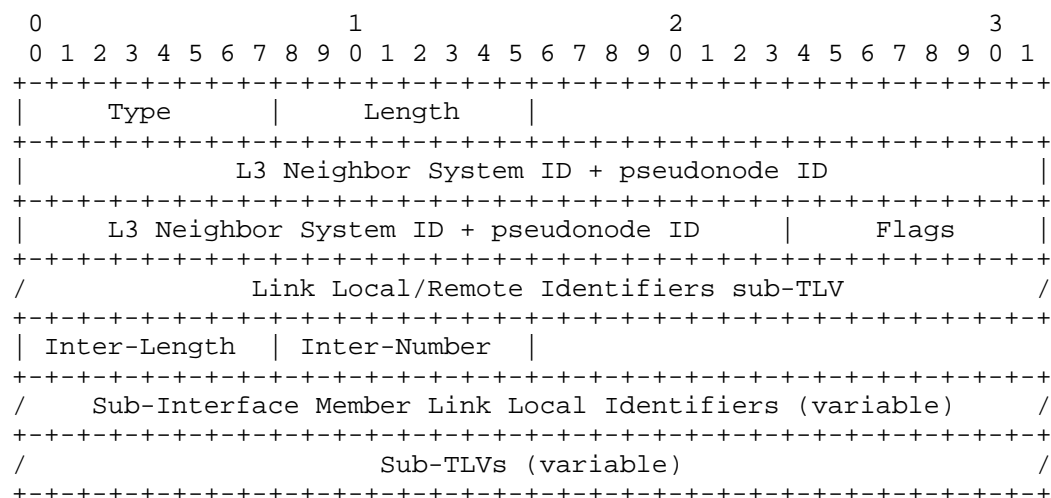


Figure 1: Format of physical local link relationship information TLV

where:

Type: TBD1, to be allocated by the IANA.

Length: the size of the value field in octets.

L3 Neighbor System ID + pseudonode ID: 7-octet length, indicates a neighbor node;

Flags: 8-bit length, currently not defined, reserved for future.

Link Local/Remote Identifiers sub-TLV: as defined in [RFC5307], it is used to identify the physical interface.

Inter-Lengh: 1-octet length, indicates the length of the following fields, including Sub-Interface Member Link Local Identifiers and Inter-Number.

Inter-Number: 1-octet length, indicates the number of the sub-interfaces.

Sub-Interface Member Link Local Identifiers: 4 * Inter-Number length, includes one or more link local identifiers for sub-interfaces.

Sub-TLVs: variable length, the Maximum Link Bandwidth Sub-TLV [RFC5305] and Unidirectional Utilized Bandwidth Sub-TLV [RFC8570] SHOULD be included in this TLV.

4. Use cases

As shown in the following figure, there are four nodes A, B, C, and D on the network. There are two links between node D and node C. IS-IS neighbor relationships are established between sub-interfaces d1 and d2 of node D and sub-interfaces c1 and c2 of node C. Sub-interfaces d1 and d2 correspond to the same physical interface d, and sub-interfaces c1 and c2 correspond to the same physical interface c. The physical bandwidth of the link between A and B is 100 Gbit/s, and 70 Gbit/s is used. The physical bandwidth of the link between C and D is 50 Gbit/s, and 20 Gbit/s is used.




Figure 3: An example for using the Sub-interface Link as Traffic Optimization Path

After the link state synchronization, node A receives and parses the related TLVs, and establishes the following mapping between the physical interface and sub-interfaces:

* Node D

- Physical interface d
 - o Bandwidth attributes of interface d (physical bandwidth and bandwidth utilization ratio)
 - o sub-interface d1
 - o sub-interface d2

* Node C

- Physical interface c
 - o Bandwidth attributes of interface c (physical bandwidth and bandwidth utilization ratio)
 - o sub-interface c1
 - o sub-interface c2

Assume that node A is the source node and node C is the destination node, the traffic rate is 60Gbit/s. In normal cases, the primary path is A->B->C. Node A calculates an alternative path A->D->C for node C, where two links are established between node D and node C through two sub-interfaces.

Node A perceives that there is a congestion on the path A->B->C, it decides to use the alternative paths for load balancing.

Based on the above established mapping relationship, node A knows that the two links corresponds to one physical interface, then uses them as a single load balancing link. Then node A will allocate 30Gbit/s on each path.

However, if node A does not know the mapping relationship, it will consider A->D->C as two independent paths, allocating 20Gbit/s on each path. This result in the traffic on the physical interfaces d exceeds its maximum bandwidth, therefore a large number of packets will be lost.

5. IANA Considerations

TBD

6. Security Considerations

TBD

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/rfc/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/rfc/rfc8174>>.

7.2. 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/rfc/rfc5305>>.

- [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/rfc/rfc8750>>.
- [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/rfc/rfc7471>>.
- [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/rfc/rfc5307>>.
- [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/rfc/rfc8570>>.

Acknowledgements

TBD

Authors' Addresses

Li Zhang
Huawei
China
Email: zhangli344@huawei.com

Jie Dong
Huawei
China
Email: jie.dong@huawei.com

Chenxi Li
Huawei
China
Email: lichenxil@huawei.com