

Link State Routing Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 2, 2021

Y. Wang
T. Zhou
Z. Hu
Huawei
October 29, 2020

IGP Extensions for Advertising Hop-by-Hop Options Header Processing
Action
draft-wang-lsr-hbh-process-00

Abstract

This document extends Node and Link attribute TLVs to Interior Gateway Protocols (IGP) to advertise the Hop-by-Hop Options header processing action and supported services (e.g. IOAM Trace Option and Alternate Marking) at node and link granularity. Such advertisements allow entities (e.g., centralized controllers) to determine whether the Hop-by-Hop Options header and specific services can be supported in a given network.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

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 May 2, 2021.

Copyright Notice

Copyright (c) 2020 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Hop-by-Hop Options Header Processing Action	4
4. Signaling Processing Action Using IS-IS	5
4.1. IS-IS Node Processing-Action Sub-TLV	5
4.2. IS-IS Link Processing-Action Sub-TLV	6
5. Signaling Processing Action Using OSPF	6
5.1. OSPF Node Processing-Action TLV	7
5.2. OSPFv2 Link Processing-Action sub-TLV	7
5.3. OSPFv3 Link Processing-Action Sub-TLV	8
6. IANA Considerations	8
7. Security Considerations	9
8. Acknowledgements	9
9. References	9
9.1. Normative References	9
9.2. Informative References	10
Authors' Addresses	10

1. Introduction

[RFC8200] specifies IPv6 extension headers, including Hop-by-Hop Options header, Destination Options header, Routing header, etc. An IPv6 packet may carry zero, one, or more extension headers that must be processed strictly in the order they appear in the packet. Except for the Hop-by-Hop Options header, other extension headers are not processed, inserted, or deleted by any transit node along a packet's delivery path, until the packet arrives at the Destination node.

As specified in [RFC8200], although the Hop-by-Hop Options header is not inserted or deleted by any transit node along a packet's delivery path, it is only examined and processed by nodes along a packet's

delivery path if they are explicitly configured to process. Besides, nodes may be configured to ignore the Hop-by-Hop Options header, drop packets containing a Hop-by-Hop Options header, or assign packets containing a Hop-by-Hop Options header to a slow processing path. In the results, devices of different vendors can be configured to process the Hop-by-Hop Options header in different ways.

Until now, the Hop-by-Hop Options header has been widely used. In-situ Operations, Administration, and Maintenance (IOAM) data fields are encapsulated in two types of extension headers in IPv6 packets, either Hop-by-Hop Options header or Destination Options header, depending on IOAM usage [I-D.ietf-ippm-ioam-ipv6-options]. For example, IOAM-tracing options are represented as an IPv6 options in Hop-by-Hop extension header. Similarly, the Alternate Marking technique can be carried by the Hop-by-Hop Options header and the Destination Options header [I-D.ietf-6man-ipv6-alt-mark]. If nodes are not explicitly configured to process the Hop-by-Hop Option header, they should ignore them. In this case, the performance measurement does not account for all links and nodes along a path. Therefore, such advertisement can be useful for entities (e.g., the centralized controller) to determine whether a specific service can be implemented in IPv6 network by encoding in the Hop-by-Hop Options header.

BGP-LS defines a way to advertise topology and associated attributes and capabilities of the nodes in that topology to a centralized controller [RFC7752]. Typically, BGP-LS is configured on a small number of nodes that do not necessarily act as head-ends. In order for BGP-LS to signal the processing action of the Hop-by-Hop Options header for all the devices in the network, the processing action SHOULD be advertised by every IGP router in the network.

This document defines a mechanism to signal the configured processing action of the Hop-by-Hop Options header and supported services at node and/or link granularity using IS-IS, OSPFv2 and OSPFv3.

2. Terminology

Following are abbreviations used in this document:

- o IGP: Interior Gateway Protocols
- o IS-IS: Intermediate System to Intermediate System
- o OSPF: Open Shortest Path First
- o BGP-LS: Border Gateway Protocol - Link State

- o NLRI: Network Layer Reachability Information

3. Hop-by-Hop Options Header Processing Action

This document defines the information of processing action formed of a tuple of a 1-octet Extension Header Options identifier and 8-bit Processing Action Flag.

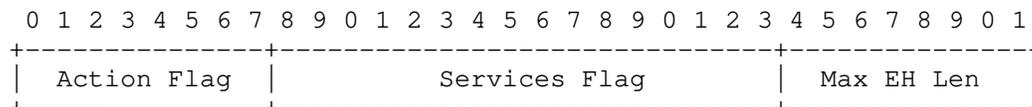


Fig. 1 Processing Action Format

Where:

- o Action Flag: A 8-bit field. The highest-order 3-bit indicates the processing action, i.e., 000 - drop packets; 001 - dispatch to control plane; 010 - forward, skip to Next Header; 011 - forward, ignoring all extension Options header; 100 - examine and process.
- o Max EH Len: A one octet field. The maximum length of the Extension Header in 8-octet units can be examined and processed at node or link granularity. The definition is same as the Next Header Length in [RFC8200].
- o Services Flag: A 16-bit bitmap. The format is defined as follows.

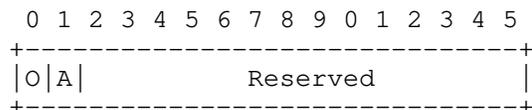


Fig. 2 Services Flag Format

where fields are defined as the following:

- o O (IOAM Trace Option) is a one-bit flag. The O flag is set to 1 if the IOAM Trace Option is supported at node or link granularity.
- o A (Alternate Marking) is a one-bit flag. The A flag is set to 1 if the Alternate Marking method is supported at node or link granularity.
- o R - reserved bits for future use. These flags MUST be zeroed on transmit and ignored on receipt.

In this document, the processing action at link granularity is defined as the supported the Hop-by-Hop Options header processing action of the interface associated with the link. When all interfaces associated with links support the same processing action, the processing action at node granularity SHOULD represent the Link processing action. Both of Node and Link processing action information are formed of a tuple of a 1-octet Extension Header Options identifier and 8-bit Processing Action Flag.

When both of Node and Link processing action are advertised, the Link processing action information MUST take precedence over the Node processing action. Besides, when a Link processing action is not signaled, then the Node processing action SHOULD be considered to be the processing action for this link.

4. Signaling Processing Action Using IS-IS

The IS-IS Extensions for advertising Router Information TLV named IS-IS Router CAPABILITY TLV [RFC7981], which allows a router to announce its capabilities within an IS-IS level or the entire routing domain.

4.1. IS-IS Node Processing-Action Sub-TLV

According to the format of IS-IS Router CAPABILITY TLV [RFC7981], the Node Processing-Action sub-TLV within the body of the IS-IS router CAPABILITY TLV is composed of three fields, a one-octet Type field, a one-octet Length field, and a 4-octet Value field. The following format is used:

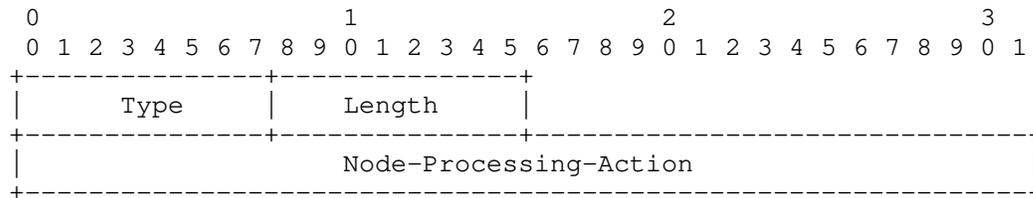


Fig. 3 IS-IS Node Processing-Action Sub-TLV Format

Where:

- o Type: To be assigned by IANA
- o Length: A one-octet field that indicates the length of the value portion in octets.
- o Node-Processing-Action: A 4-octet field, which is same as defined in Section 3.

4.2. IS-IS Link Processing-Action Sub-TLV

The Link Processing-Action sub-TLV is defined for TLVs 22, 23, 25, 141, 222, and 223 to carry the Processing-Action information of the interface associated with the link. The Link Processing-Action sub-TLV is formed of three fields, a one-octet Type field, a one-octet Length field, and a 4-octet Value field. The following format is used:

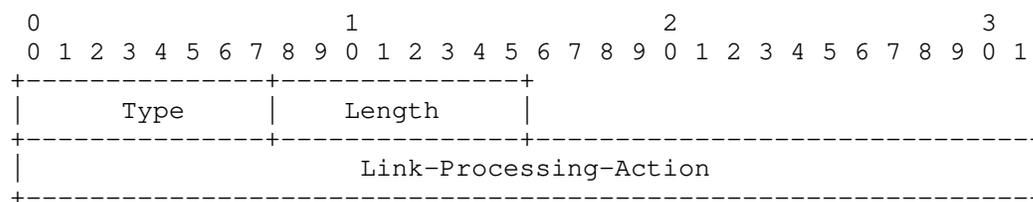


Fig. 4 IS-IS Link Processing-Action Sub-TLV Format

Where:

- o Type: To be assigned by IANA
- o Length: A one-octet field that indicates the length of the value portion in octets.
- o Link-Processing-Action: A 4-octet field, which is same as defined in Section 3.

5. Signaling Processing Action Using OSPF

Given that OSPF uses the options field in LSAs and hello packets to advertise optional router capabilities [RFC7770], this document defines a new Node Processing-Action TLV within the body of the OSPF RI Opaque LSA [RFC7770] to carry the Processing Action of the router originating the RI LSA.

This document defines the Link Processing-Action sub-TLV to carry the Processing-Action information of the interface associated with the link. For OSPFv2, the link-level Processing-Action information is advertised as a sub-TLV of the OSPFv2 Extended Link TLV as defined in [RFC7684]. For OSPFv3, the link-level Processing-Action information is advertised a sub-TLV of the E-Router-LSA TLV as defined in [RFC8362].

5.1. OSPF Node Processing-Action TLV

The Node Processing-Action TLV is composed of three fields, a 2-octet Type field, a 2-octet Length field, and a 4-octet Value field. The following format is used:

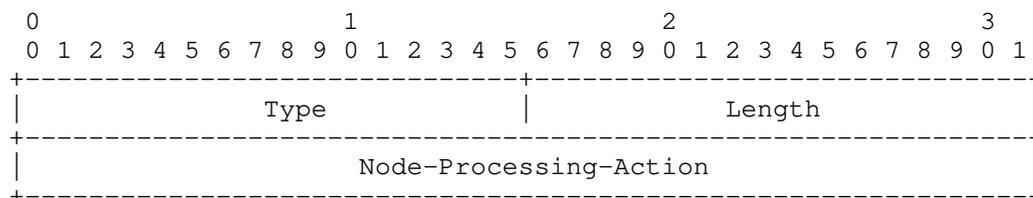


Fig. 5 OSPF Node Processing-Action TLV

Where:

- o Type: To be assigned by IANA
- o Length: A 2-octet field that indicates the length of the value field.
- o Node-Processing-Action: A 4-octet field, which is as defined in Section 3.

5.2. OSPFv2 Link Processing-Action sub-TLV

The Link Processing-Action sub-TLV encoded in the OSPFv2 Extended Link TLV as defined in [RFC7684], which is constructed of three fields, a 2-octet Type field, a 2-octet Length field, and a 4-octet Value field. The following format is used:

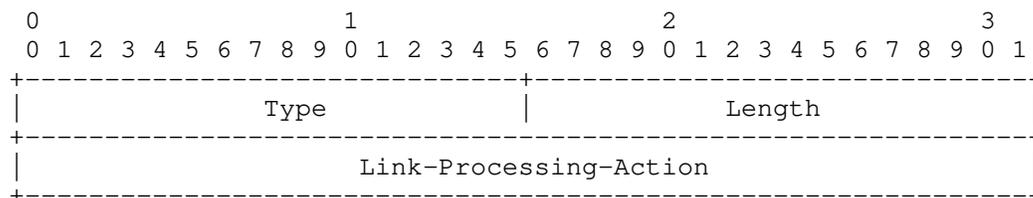


Fig. 6 OSPFv2 Link Processing-Action sub-TLV

Where:

- o Type: To be assigned by IANA

- o Length: A 2-octet field that indicates the length of the value field.
- o Link-Processing-Action: A 4-octet field, which is as defined in Section 3.

5.3. OSPFv3 Link Processing-Action Sub-TLV

The Link Processing-Action sub-TLV encoded in the OSPFv3 E-Router-LSA TLV as defined in [RFC8362], which is constructed of three fields, a 2-octet Type field, a 2-octet Length field, and a 4-octet Value field. The following format is used:

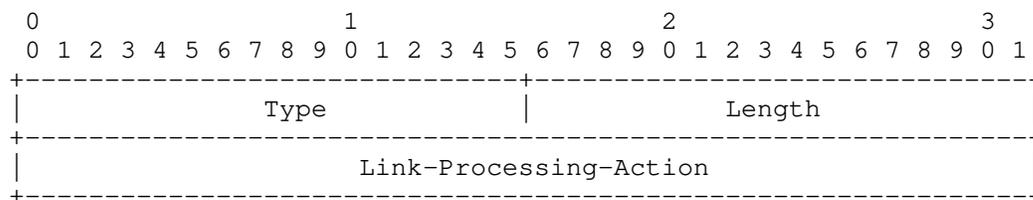


Fig. 7 OSPFv3 Link Processing-Action sub-TLV

Where:

- o Type: To be assigned by IANA
- o Length: A 2-octet field that indicates the length of the value field.
- o Link-Processing-Action: A 4-octet field, which is as defined in Section 3.

6. IANA Considerations

IANA is requested to allocate values for the following new TLV and sub-TLVs.

Type	Description
TBA	IS-IS Node Processing-Action Sub-TLV
TBA	IS-IS Link Processing-Action Sub-TLV

Type	Description
TBA	OSPF Node Processing-Action TLV
TBA	OSPFv2 Link Processing-Action sub-TLV
TBA	OSPFv3 Link Processing-Action sub-TLV

7. Security Considerations

This document introduces new IGP Node and Link Attribute TLVs and sub-TLVs for advertising processing actions of the Hop-by-Hop Options header at node and/or link granularity. It does not introduce any new security risks to IS-IS, OSPFv2 and OSPFv3.

8. Acknowledgements

TBD

9. References

9.1. 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>>.
- [RFC7684] "OSPFv2 Prefix/Link Attribute Advertisement", <<https://www.rfc-editor.org/info/rfc7684>>.
- [RFC7752] "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", <<https://datatracker.ietf.org/doc/rfc7752/>>.
- [RFC7770] "Extensions to OSPF for Advertising Optional Router Capabilities", <<https://www.rfc-editor.org/info/rfc7770>>.
- [RFC7981] "IS-IS Extensions for Advertising Router Information", <<https://www.rfc-editor.org/info/rfc7981>>.
- [RFC8200] "Internet Protocol, Version 6 (IPv6) Specification", <<https://datatracker.ietf.org/doc/rfc8200/>>.
- [RFC8362] "OSPFv3 Link State Advertisement (LSA) Extensibility", <<https://www.rfc-editor.org/info/rfc8362>>.

9.2. Informative References

[I-D.ietf-6man-ipv6-alt-mark]
"IPv6 Application of the Alternate Marking Method",
<<https://datatracker.ietf.org/doc/draft-ietf-6man-ipv6-alt-mark/>>.

[I-D.ietf-ippm-ioam-ipv6-options]
"In-situ OAM IPv6 Options",
<<https://datatracker.ietf.org/doc/draft-ietf-ippm-ioam-ipv6-options/>>.

Authors' Addresses

Yali Wang
Huawei
156 Beiqing Rd., Haidian District
Beijing
China

Email: wangyalil1@huawei.com

Tianran Zhou
Huawei
156 Beiqing Rd., Haidian District
Beijing
China

Email: zhoutianran@huawei.com

Zhibo Hu
Huawei
156 Beiqing Rd., Haidian District
Beijing
China

Email: huzhibo@huawei.com