



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. Requirements Language](#)
- [3. Terminology](#)
- [4. Background](#)
- [5. Hop-by-Hop Header Processing Procedures](#)
  - [5.1. Hop-by-Hop Options Per Packet](#)
  - [5.2. Hop-by-Hop Headers Processing](#)
  - [5.3. Configuration](#)
- [6. Hop-by-Hop Option Header Alignment](#)
  - [6.1. Hop-by-Hop Options that Meet Alignment Requirement](#)
  - [6.2. Hop-by-Hop Options that Must be Deprecated or Modified](#)
  - [6.3. Destination Options and Deprecated Hop-by-Hop Options](#)
- [7. New Hop-by-Hop Options](#)
- [8. IANA Considerations](#)
- [9. Security Considerations](#)
- [10. Acknowledgments](#)
- [11. Change log \[RFC Editor: Please remove\]](#)
- [12. Normative References](#)
- [13. Informative References](#)
- [Authors' Addresses](#)

## 1. Introduction

This document specifies procedures for how IPv6 Hop-by-Hop options are processed. It modifies the procedures specified in the IPv6 Protocol Specification (RFC8200) to make processing in routers more practical with the goal of making IPv6 Hop-by-Hop options useful to deploy in the Internet.

When published this document updates [[RFC8200](#)].

The current list of defined Hop-by-Hop options can be found at [[IANA-HBH](#)].

## 2. 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.

### 3. Terminology

This document uses the following loosely defined terms:

\*Fast path: Hardware or Application-Specific Integrated Circuit (ASIC) processing path for packets. This is the usual processing path within a router for IP datagrams. It is sometimes also described as the forwarding plane.

\*Slow path: Software processing path for packets. This is a special processing path for packets that require special processing or differ from assumptions made in fast-path heuristics, or for router control protocols. It is sometimes also described as the control plane.

### 4. Background

In the first version of the IPv6 specification, Hop-by-Hop options were required to be processed by all nodes, routers and hosts. This proved to not be practical in high speed routers due to several factors, including:

\*Inability to process the hop-by-hop options at wire speed in hardware.

\*Hop-by-Hop options would be sent to the "slow path". This could degrade the routers performance and it's ability to process important control traffic.

\*A mechanism that forces external packets to the routers "slow path" could be exploited as a Denial of Service attack on the router.

\*Packets could contain multiple Hop-by-Hop options making the previous issues worse by increasing the complexity required to process them.

When the IPv6 Specification was updated and published in July 2017 as [[RFC8200](#)], the procedures relating to hop-by-hop options are as follows:

Extension headers (except for the Hop-by-Hop Options header) are not processed, inserted, or deleted by any node along a packet's delivery path, until the packet reaches the node (or each of the set of nodes, in the case of multicast) identified in the Destination Address field of the IPv6 header.

The Hop-by-Hop Options header is not inserted or deleted, but may be examined or processed by any node along a packet's delivery path, until the packet reaches the node (or each of the set of

nodes, in the case of multicast) identified in the Destination Address field of the IPv6 header. The Hop-by-Hop Options header, when present, must immediately follow the IPv6 header. Its presence is indicated by the value zero in the Next Header field of the IPv6 header.

NOTE: While [RFC2460] required that all nodes must examine and process the Hop-by-Hop Options header, it is now expected that nodes along a packet's delivery path only examine and process the Hop-by-Hop Options header if explicitly configured to do so.

The changes meant that an implementation complied with the IPv6 specification even if it did not process hop-by-hop options, and that it was expected that routers would add configuration information to control which hop-by-hop options they would process.

Unfortunately, this did not improve the processing of Hop-by-Hop options and they are often not possible to be deployed and used in the Internet. It merely documented how they were being used in the Internet.

In hindsight, hop-by-hop options were still not practical to be used widely in the Internet. Many operational routers are configured to drop all packets containing a hop-by-hop option header. The main issues remain:

- \*Routers have been configured to drop transit data packets that would be processed in the fast path, this includes dropping hop-by-hop packets that would need to be processed in the fast path.

- \*Allowing multiple hop-by-hop options in a single packet makes it even more expensive in router resources to process these packets in the fast path. It adds complexity to the number of permutations that might need to be processed.

- \*Any mechanism that can be used externally to force packets into the router's slow path can be exploited as a denial of service attack on a transit router by saturating the resources need for router management protocols (e.g., routing protocols, network management protocols, etc.) that may cause the router to fail. This issue for the Router Alert option, which intentionally places packets on the slow path, is discussed in [[RFC6398](#)]. Section 3 of this RFC includes a good summary:

"In a nutshell, the IP Router Alert Option does not provide a convenient universal mechanism to accurately and reliably distinguish between IP Router Alert packets of interest and unwanted IP Router Alert packets. This, in turn, creates a security concern when the IP Router Alert Option is used, because, short of appropriate router-implementation-specific

mechanisms, the router slow path is at risk of being flooded by unwanted traffic."

There has been some academic research that discussed the general problem with dropping packets containing IPv6 extension headers, including the Hop-by-Hop Options header. For example [\[Hendriks\]](#) states that "dropping all packets with Extension Headers, is a bad practice", and that "The share of traffic containing more than one EH however, is very small. For the design of hardware able to handle the dynamic nature of EHs, we therefore recommend to support at least one EH"

This document defines a set of procedures for the hop-by-hop option header that make the processing of hop-by-hop options practical in modern transit routers.

## **5. Hop-by-Hop Header Processing Procedures**

This section describes three sets of changes to [\[RFC8200\]](#).

### **5.1. Hop-by-Hop Options Per Packet**

The Hop-by-Hop Option Header as defined in Section 4.3 of [\[RFC8200\]](#) is identified by a Next Header value of 0 in the IPv6 header. Section 4.1 requires a Hop-by-Hop Options header to appear immediately after the IPv6 header. [\[RFC8200\]](#) also requires that a Hop-by-Hop Options header can only appear once in a packet.

The Hop-by-Hop Options Header as defined in [\[RFC8200\]](#) can contain one or more Hop-by-Hop options. This document updates [\[RFC8200\]](#) such that there MUST only be one option contained in an Hop-by-Hop Options header in a single packet. The motivation for this change is to simplify the processing of Hop-by-Hop options in the fast path.

Nodes creating packets with a Hop-by-Hop option headers MUST only include a single option in the packet. This also requires that all Hop-by-Hop Options MUST be 8-octet aligned. See [Section 6](#) for more detail.

Nodes processing a packet with a Hop-by-Hop options header MUST discard the packet if there is more than one Hop-by-Hop option in that header.

### **5.2. Hop-by-Hop Headers Processing**

Section 4.2 of [\[RFC8200\]](#) defines an Option Type field in options that controls how the option is processed if the IPv6 node processing the option header does not recognize the Option Type. This document modifies that behavior for options carried in an Hop-by-Hop Option header in two ways.

IPv6 nodes MUST only process an Hop-by-Hop Options header if it can be done in the fast path of the router. If the IPv6 node can not process the Hop-by-Hop Option header in the fast path it MUST skip over this option and continue processing the header normally.

Section 4.2 of [[RFC8200](#)] defines the Option Type identifiers as internally encoded such that their highest-order 2 bits specify the action that must be taken if the processing IPv6 node does not recognize the Option Type. This document modifies this behaviour for Hop-by-Hop options as follows:

00 - skip over this option and continue processing the header.

01 - discard the packet.

10 - discard the packet.

11 - discard the packet.

The motivation for this change is to simplify what the router needs to do in the fast path when it can not recognize the Option Type. It is no longer required to send ICMP Parameter Problem messages.

### 5.3. Configuration

Section 4 of [[RFC8200](#)] allows for a router to control it's processing of IPv6 Hop-by-Hop options by local configuration. The text is:

NOTE: While [[RFC2460](#)] required that all nodes must examine and process the Hop-by-Hop Options header, it is now expected that nodes along a packet's delivery path only examine and process the Hop-by-Hop Options header if explicitly configured to do so.

A possible approach to implementing this is to maintain a lookup table based on Option Type of the IPv6 options that are supported in the fast path. This would allow for a node to quickly determine if an option is supported and can be processed. If the option is not supported, then the node processes it as described in [Section 5.2](#) of this document.

The actions of the lookup table SHOULD be configurable by the operator of the router.

## 6. Hop-by-Hop Option Header Alignment

[[RFC8200](#)] requires all extension headers to be 8-octet aligned. The text from Section 4 is:

Each extension header is an integer multiple of 8 octets long, in order to retain 8-octet alignment for subsequent headers.

Two types of Pad options are defined to keep this alignment, that is Pad1 and PadN, if the options in an Destination Option Header or Hop-by-Hop Option Header do not result in 8-octet alignment.

This document requires all Hop-by-Hop Options to be 8-octet aligned (see [Section 5.1](#)). This eliminates the need for Pad options to be used in the Hop-by-Hop Option header to make them 8-octet aligned.

The current list of defined Destination and Hop-by-Hop options can be found at [[IANA-HBH](#)].

The following sections list current Hop-by-Hop Options that meet the alignment requirement defined here ([Section 6.1](#)) or if they need to be deprecated or modified ([Section 6.2](#)). Hop-by-Hop Options that have been deprecated, and Destination Options are listed separately in [Section 6.3](#). Pad Options (Pad1 and PadN), and RFC3692-style Experiment options are not listed.

### 6.1. Hop-by-Hop Options that Meet Alignment Requirement

The following Hop-by-Hop Options meet the alignment requirements defined in this section:

- \*Jumbo Payload [[RFC2675](#)]
- \*Path MTU Record Option [[I-D.ietf-6man-mtu-option](#)]
- \*RPL Option [[I-D.ietf-roll-useofrplinfo](#)] \*\*\*
- \*Quick-Start [[RFC4782](#)]
- \*CALIPSO [[RFC5570](#)] \*\*\*
- \*SMF\_DPD [[RFC6621](#)] \*\*\*
- \*ILNP Nonce [[RFC6744](#)] \*\*\*
- \*MPL Option [[RFC7731](#)]

\*\*\* Note: Need to verify that these options are 8-octet aligned.

### 6.2. Hop-by-Hop Options that Must be Deprecated or Modified

The Hop-by-Hop Options listed in this section do not meet the alignment requirements defined in this section. They either need to be deprecated or modified to support 8-octet alignment without the

use of Pad options. If there is interest in modifying them, it should be straightforward to make them have 8-octet alignment.

- \*Router Alert [[RFC2711](#)]
- \*IP\_DFF [[RFC6971](#)]
- \*IOAM [[I-D.ietf-ippm-ioam-ipv6-options](#)]

### 6.3. Destination Options and Deprecated Hop-by-Hop Options

The options listed in this section are either Destination Options or Deprecated Hop-by-Hop Options. 8-octet alignment is not an issue.

- \*Tunnel Encapsulation Limit [[RFC2473](#)]
- \*RPL Option (DEPRECATED) [[RFC6553](#)]
- \*Endpoint Identification (DEPRECATED) [CHARLES LYNN]
- \*MPL Option (DEPRECATED) [[RFC7731](#)]
- \*Home Address [[RFC6275](#)]
- \*Line-Identification Option [[RFC6788](#)]
- \*Performance and Diagnostic Metrics (PDM) [[RFC8250](#)]

### 7. New Hop-by-Hop Options

Any new IPv6 Hop-by-Hop option defined in the future should have 8-octet alignment without the use of any Pad options. This requirement modifies Section 4.8 of [[RFC8200](#)].

### 8. IANA Considerations

This draft has no actions for the IANA.

### 9. Security Considerations

TBD

### 10. Acknowledgments

Helpful comments were received from Brian Carpenter, Ron Bonica, Ole Troan, [your name here], and other members of the 6MAN working group.

### 11. Change log [RFC Editor: Please remove]

draft-hinden-6man-hbh-processing-00, 2020-Nov-29: Initial draft.

### 12. Normative References

[IANA-HBH] "Destination Options and Hop-by-Hop Options", <<https://www.iana.org/assignments/ipv6-parameters/ipv6-parameters.xhtml#ipv6-parameters-2>>.



**[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>>.

**[RFC8200]**

Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

### 13. Informative References

**[Hendriks]**

Hendriks, L., Velan, P., Schmidt, R.O., Boer, P., and A. Aiko, "Threats and Surprises behind IPv6 Extension Headers", , , August 2017, <[http://dl.ifip.org/db/conf/tma/tma2017/tma2017\\_paper22.pdf](http://dl.ifip.org/db/conf/tma/tma2017/tma2017_paper22.pdf)>.

**[I-D.ietf-6man-mtu-option]**

Hinden, R. and G. Fairhurst, "IPv6 Minimum Path MTU Hop-by-Hop Option", Work in Progress, Internet-Draft, draft-ietf-6man-mtu-option-04, 23 October 2020, <<https://tools.ietf.org/html/draft-ietf-6man-mtu-option-04>>.

**[I-D.ietf-ippm-ioam-ipv6-options]**

Bhandari, S., Brockners, F., Pignataro, C., Gredler, H., Leddy, J., Youell, S., Mizrahi, T., Kfir, A., Gafni, B., Lapukhov, P., Spiegel, M., Krishnan, S., Asati, R., and M. Smith, "In-situ OAM IPv6 Options", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-ipv6-options-04, 1 November 2020, <<https://tools.ietf.org/html/draft-ietf-ippm-ioam-ipv6-options-04>>.

**[I-D.ietf-roll-useofrplinfo]**

Robles, I., Richardson, M., and P. Thubert, "Using RPI Option Type, Routing Header for Source Routes and IPv6-in-IPv6 encapsulation in the RPL Data Plane", Work in Progress, Internet-Draft, draft-ietf-roll-

useofrplinfo-42, 12 November 2020, <<https://tools.ietf.org/html/draft-ietf-roll-useofrplinfo-42>>.

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473, December 1998, <<https://www.rfc-editor.org/info/rfc2473>>.

[RFC2675] Borman, D., Deering, S., and R. Hinden, "IPv6 Jumbograms", RFC 2675, DOI 10.17487/RFC2675, August 1999, <<https://www.rfc-editor.org/info/rfc2675>>.

[RFC2711] Partridge, C. and A. Jackson, "IPv6 Router Alert Option", RFC 2711, DOI 10.17487/RFC2711, October 1999, <<https://www.rfc-editor.org/info/rfc2711>>.

[RFC4782] Floyd, S., Allman, M., Jain, A., and P. Sarolahti, "Quick-Start for TCP and IP", RFC 4782, DOI 10.17487/RFC4782, January 2007, <<https://www.rfc-editor.org/info/rfc4782>>.

[RFC5570] StJohns, M., Atkinson, R., and G. Thomas, "Common Architecture Label IPv6 Security Option (CALIPSO)", RFC 5570, DOI 10.17487/RFC5570, July 2009, <<https://www.rfc-editor.org/info/rfc5570>>.

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.

[RFC6398] Le Faucheur, F., Ed., "IP Router Alert Considerations and Usage", BCP 168, RFC 6398, DOI 10.17487/RFC6398, October 2011, <<https://www.rfc-editor.org/info/rfc6398>>.

[RFC6553] Hui, J. and JP. Vasseur, "The Routing Protocol for Low-Power and Lossy Networks (RPL) Option for Carrying RPL Information in Data-Plane Datagrams", RFC 6553, DOI 10.17487/RFC6553, March 2012, <<https://www.rfc-editor.org/info/rfc6553>>.

[RFC6621] Macker, J., Ed., "Simplified Multicast Forwarding", RFC 6621, DOI 10.17487/RFC6621, May 2012, <<https://www.rfc-editor.org/info/rfc6621>>.

[RFC6744] Atkinson, RJ. and SN. Bhatti, "IPv6 Nonce Destination Option for the Identifier-Locator Network Protocol for IPv6 (ILNPv6)", RFC 6744, DOI 10.17487/RFC6744, November 2012, <<https://www.rfc-editor.org/info/rfc6744>>.

[RFC6788] Krishnan, S., Kavanagh, A., Varga, B., Ooghe, S., and E. Nordmark, "The Line-Identification Option", RFC 6788, DOI

10.17487/RFC6788, November 2012, <<https://www.rfc-editor.org/info/rfc6788>>.

[RFC6971] Herberg, U., Ed., Cardenas, A., Iwao, T., Dow, M., and S. Cespedes, "Depth-First Forwarding (DFF) in Unreliable Networks", RFC 6971, DOI 10.17487/RFC6971, June 2013, <<https://www.rfc-editor.org/info/rfc6971>>.

[RFC7731] Hui, J. and R. Kelsey, "Multicast Protocol for Low-Power and Lossy Networks (MPL)", RFC 7731, DOI 10.17487/RFC7731, February 2016, <<https://www.rfc-editor.org/info/rfc7731>>.

[RFC8250] Elkins, N., Hamilton, R., and M. Ackermann, "IPv6 Performance and Diagnostic Metrics (PDM) Destination Option", RFC 8250, DOI 10.17487/RFC8250, September 2017, <<https://www.rfc-editor.org/info/rfc8250>>.

#### Authors' Addresses

Robert M. Hinden  
Check Point Software  
959 Skyway Road  
San Carlos, CA 94070  
United States of America

Email: [bob.hinden@gmail.com](mailto:bob.hinden@gmail.com)

Godred Fairhurst  
University of Aberdeen  
School of Engineering  
Fraser Noble Building  
Aberdeen  
AB24 3UE  
United Kingdom

Email: [gorry@erg.abdn.ac.uk](mailto:gorry@erg.abdn.ac.uk)