

6MAN Working Group
Internet-Draft
Intended status: Standards Track
Expires: 10 December 2025

D. Eastlake
Independent
J. Dong
Huawei Technologies
8 June 2025

Transient Hiding of Hop-by-Hop Options
draft-eastlake-6man-hide-options-09

Abstract

There are an increasing number of IPv6 hop-by-hop options specified but such IPv6 options are poorly handled, particularly by high-speed routers in the core Internet where packets having options may be discarded. This document proposes a simple method of transiently hiding such options for part of a packet's path to protect the packet from discard or mishandling.

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 10 December 2025.

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
1.1. Conventions Used in This Document	2
2. IP Options and Option Handling Problems	3
2.1. IPv6 Options	4
3. Overview of a Solution	5
3.1. Transiently Hiding IPv6 Options	7
3.2. Evolution to Greater Option Support	7
4. IANA Considerations	7
5. Security Considerations	8
6. Normative References	8
7. Informative References	8
Appendix A. Revision History	9
A.1. -00 to -01	9
A.2. -01 to -02	9
A.3. -02 to -03 to -04	9
A.4. -04 to -05	9
A.5. -05 to -06 to -07	9
Acknowledgments	10
Authors' Addresses	10

1. Introduction

As discussed in [Options3] there are an increasing number of IPv6 hop-by-hop options being specified. But such IPv6 options, are poorly handled, particularly by high-speed routers in the core Internet where packets having options may be discarded. This document proposes a simple method of transiently hiding such options for part of a packet's path to protect the packet from discard or mishandling.

1.1. Conventions Used in This Document

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.

Terms:

ASIC - Application Specific Integrated Circuit.

field - an area of one or more contiguous bits within a larger data structure.

2. IP Options and Option Handling Problems

This Section is informational and intended to provide background information.

In the early days of the Internet, much of the traffic was text, transmission speeds were slow, and IP routers were commonly small general-purpose computers. Under these conditions, parsing IP headers with various options or combinations of options, handling variable length options, etc., was relatively easy.

However, as time passed, the Internet increased in size, bandwidth grew including more voluminous media such as video, transmission speeds increased enormously, and latency/responsiveness requirements became more stringent. These requirements have led to IP routers, especially in the core of the Internet, becoming less flexible and more specialized. To be able to handle data faster and more efficiently, such core IP routers are divided into a forwarding plane and a control plane where the forwarding plan handles the usual data forwarding while the control plan handles routing control messages and other packets that the data plane cannot handle. In some IP routers, the forwarding plane is implemented with Application Specific Integrated Circuits (ASICs) that are inflexible and may need fields they examine in an IP packet header to be at a fixed offset from the beginning of the packet. Meanwhile, the control plane may be implemented through a general-purpose computer which can only handle a limited number of packets per unit time.

For these reasons, many IP routers do not implement many or any types of IPv6 Hop-by-Hop options (or IPv4 [RFC0791] header options) except through the control plane which has limited capacity. Sending packets with such options to the control plane can overwhelm the control plane and interfere with routing control messages or other critical functions. Very often, particularly for IP routers handling a large volume of traffic, a strategy is adopted of dropping IP packets with such header options or ignoring the header options.

See [Options3] for a further discussion of these option handling problems.

2.1. IPv6 Options

Figure 1 shows the IPv6 header [RFC8200]. The initial 4-bit Version field indicates the IP version number and would be 0b0110 to indicate the value 6.

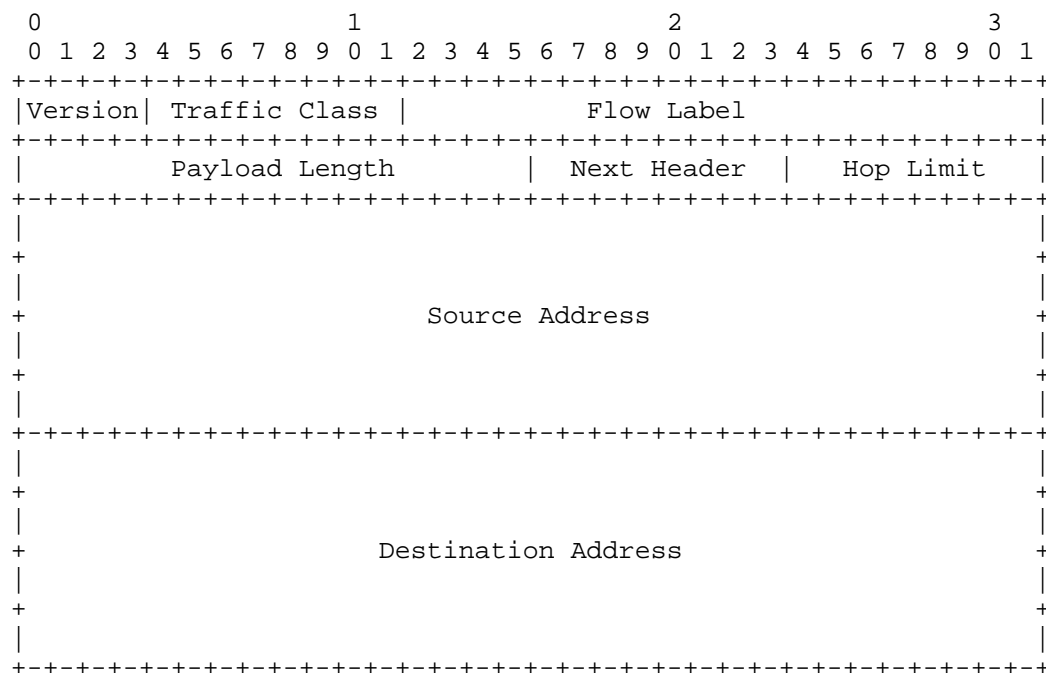


Figure 1: IPv6 Header

The value of the 8-bit Next Header field specifies the type and format of information immediately following the header. For example, a value of 17 in the Next Header field indicates that the header is immediately followed by a User Datagram Protocol (UDP) message and a value of 6 would indicate the header is followed by a Transmission Control Protocol (TCP) message. In some cases, the data immediately after the IPv6 header can be a header that itself includes a Next Header field for the type of data following it and so on as shown in Figure 2. Such headers, after the initial IPv6 header and before the main payload, are called Extension Headers and can be viewed as extensions to the IPv6 header. At this time specified extension headers include the six listed below, additional extension headers have been proposed, and likely more extension headers will be proposed and specified in the future.

Specified extension headers:

- * Hop-by-Hop Options
- * Fragment
- * Destination Options
- * Routing
- * Authentication
- * Encapsulating Security Payload

In the two "options" types of extension header, the "Hop-by-Hop Options" and "Destination Options", the extension header content is further structured into options each of which, except for a one byte "pad1" option, is an 8-bit type followed by an 8-bit option length, followed by the option value. Hop-by-Hop options were initially specified to require that every router pay attention to them. While this has been relaxed in the most recent IPv6 specification, they are still sometimes viewed as imposing a burden on every IP router through which they pass.

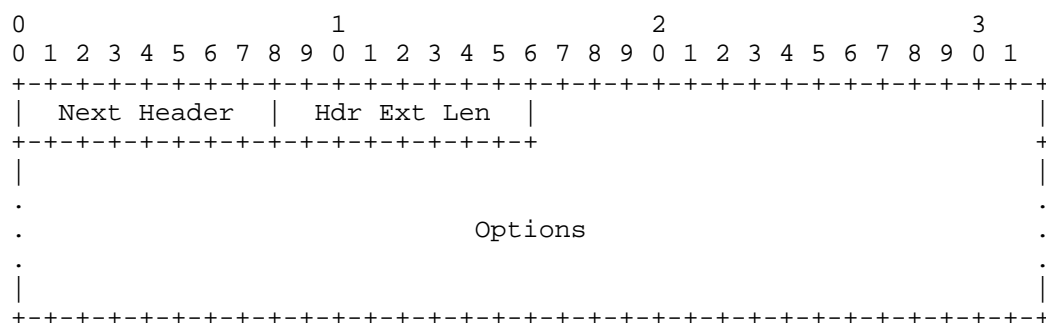


Figure 2: IPv6 Option Extension Header

3. Overview of a Solution

Figure 3 shows a simplified high-level view of a network path between two hosts within local networks through the Internet core. In reality there are likely to be more levels with a local network, whether a home, office, data center, or whatever, usually connected through one or more levels of lower tier service provider before connecting to a Tier 1 provider that connects to the Internet core, also known as the default free zone.

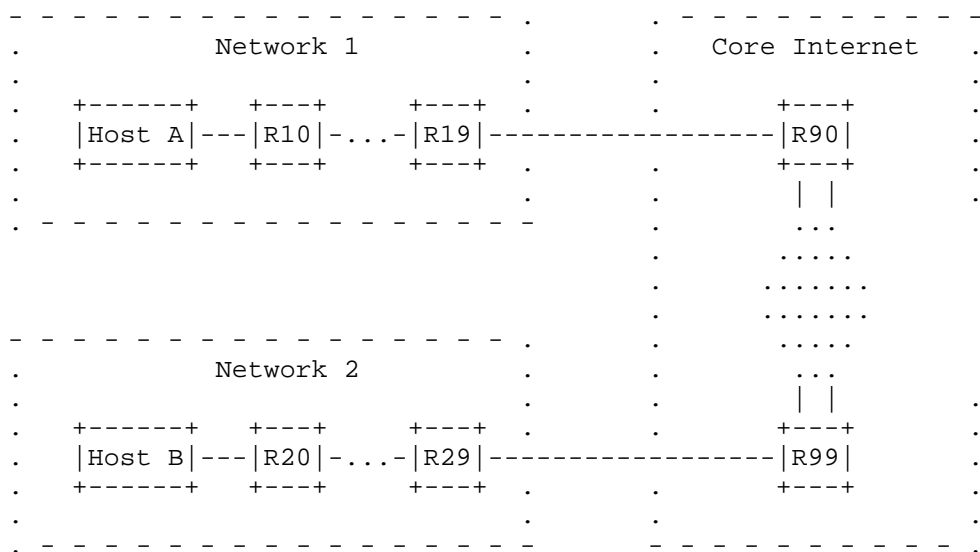


Figure 3: High Level View of an Internet Path

There are efforts to improve and streamline handling of IPv6 Hop-by-Hop options such as the methods in [Options1] and [Options2]. However, even if such a method were popular and deployed in some network areas, there is likely to be significant delay before it would be deployed in most of the Internet core. While some Internet core routers may ignore options, others discard packets with options and, as long as there is a significant chance of such discard, options are rendered essentially useless on paths through the core.

A solution is to hide options before IP packets arrive at the core. This document specifies a method so that this hiding is done in an easily detectable and reversible fashion and thus options can be easily unhidden after leaving the core. IPv6 Hop-by-Hop options so hidden might not be effective in the core but the situation can be an improvement over the traffic using such options being discarded.

This solution requires destination support but that should be knowable in many cases such as traffic between branches of the same company or between a customer and a data center.

To obtain more uniform handling of packets in a flow, it may be desirable to treat all packet in the flow as if they had such options in that the packet would be transformed to hide and unhide options even if there were none. Alternatively, this transformation could be applied to all packets starting with the first having a problematic option.

3.1. Transiently Hiding IPv6 Options

IPv6 Hop-by-Hop options are hidden by replacing the zero Next Header field in the IPv6 Header by the opaque IP protocol number TBD. This is a very simple modification of one 8-bit field in a fixed location that has no effect of the size of the packet and leaves the Traffic Class, Flow Label, Hop Limit, etc., unchanged. They are unhidden by changing this opaque IP protocol number in the IPv6 header back to zero. The points of hiding and un hiding in the packet's path (or paths if multicast) should be chosen to maximize the routers at the beginning and end of the path (Figure 3) that implement the options while minimizing the chance of unwanted packet discard or other damaging behavior.

3.2. Evolution to Greater Option Support

This solution supports the evolution of the Internet toward more widespread support or beneficent handling of hop-by-hop options as follows:

- * As acceptable option handling is more widely implemented, probably starting at lower bandwidth routers nearer the edge, the boundaries at which options are hidden and unhidden can migrate closer to the core.
- * If scattered core routers improve to provide acceptable option handling, they can recognize the opaque protocol number and perform options, perhaps in a limited way, on packets where those options are hidden to unimproved routers.

4. IANA Considerations

IANA is requested to assign a number from the "Assigned Internet Protocol Numbers" registry as follows:

Decimal	Keyword	Protocol	IPv6 Ex Hdr	Reference
TBD	Opaque	Opaque		[this document]

Table 1

5. Security Considerations

The use of the opaque IP Protocol to mask options is intended to disable normal analysis of the following packet content, specifically hop-by-hop options in the IP header. This would make firewalls, deep packet analysis, and the like less effective. However, such methods are less likely to be present in the network core where packets might be obscured. Furthermore, firewalls tend to only admit packets with known permissible values in protocol header fields such as the IP protocol field. The rejection by a firewall of a packet with the opaque IP protocol value will protect the nodes behind that firewall from possible damage due to the receipt of a packet modified as specified in this document. If the firewall does know the opaque IP Protocol value, it should be configured to treat packets with that value safely, possibly by reversing the option hiding transformation.

Performing the specified option hiding header modification only for packets with options would reveal which packets have that characteristic and might cause re-ordering of those packets with respect to other packets in the same flow. Thus, it may be beneficial to perform the specified modification on all packets in an order sensitive flow.

Should an IPv6 packet modified to hide options get through to a host that does not understand this modification, it would likely be discarded due to having an unknown IP Protocol.

[More to be added ?]

6. 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>>.
- [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>>.

7. Informative References

- [Options1] Li, Z., Peng, S., and G. Mishra, "Hop-by-Hop Forwarding Options Header", Work in progress, February 2021, <<https://datatracker.ietf.org/doc/draft-li-6man-hbh-fwd-hdr/>>.
- [Options2] Hinden, R. and G. Fairhurst, "IPv6 Hop-by-Hop Options Processing Procedures", Work in progress, 4 July 2023, <<https://datatracker.ietf.org/doc/draft-ietf-6man-hbh-processing/>>.
- [Options3] Li, Z., Peng, S., Xie, C., Qin, Z., and G. Mishra, "Operational Issues with Processing of the Hop-by-Hop Options Header", Work in progress, 10 March 2023, <<https://datatracker.ietf.org/doc/draft-ietf-v6ops-hbh/>>.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.

Appendix A. Revision History

RFC Editor: Please delete this appendix before publication.

A.1. -00 to -01

Minor editorial changes. Add more Security Considerations. Add Acknowledgements section.

A.2. -01 to -02

Delete IPv4 material. It was a bit complex and almost no one really cares about IPv4 options. Also, minor editorial changes.

A.3. -02 to -03 to -04

Minor changes.

A.4. -04 to -05

Convert to XMLv3. Add author.

A.5. -05 to -06 to -07

Minor changes.

Acknowledgments

The helpful comments of the following are gratefully acknowledged:

Peng Shuping

Authors' Addresses

Donald E. Eastlake 3rd
Independent
2386 Panoramic Circle
Apopka, Florida 32703
United States of America
Phone: +1-508-333-2270
Email: d3e3e3@gmail.com

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