INTERNET-DRAFT
Intended status: Proposed Standard
Expires: April 18, 2022

D. Eastlake
Futurewei Technologies

October 19, 2021

Transient Hiding of Hop-by-Hop Options
<draft-eastlake-6man-hide-options-01.txt>

Abstract

There are increasing requests for a variety IPv6 hop-by-hop options but such IPv6 options and all IPv4 options, are poorly handled, particularly by high-speed routers in the core Internet where packets having options are commonly 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.

Status of This Memo

Table of Contents

1. Introduction

As discussed in [Options3] there are increasing requests for a variety IPv6 hop-by-hop options but such IPv6 options and all IPv4 options, are poorly handled, particularly by high-speed routers in the core Internet where packets having options are commonly 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.

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 structure.

2. IP Options and Option Handling Problems

This Section 2 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 the Internet increased in size, bandwidth grew including more voluminous media such as video, transmission speeds increased enormously, and latency/responsiveness requirements became much more stringent. This leads 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 relatively low power general purpose computer which can only handle a small 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 header options except through the control plane which is relatively slow. 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 IPv4 header options and IPv6 Hop-by-Hop header options.

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

Further details concerning IPv6 and IPv4 options are given in the subsections below.

2.1 IPv6 Options

Figure 1 shows the IPv6 header [RFC8200]. The value of the initial 4-bit Version field indicates the IP version number and has the value 6.

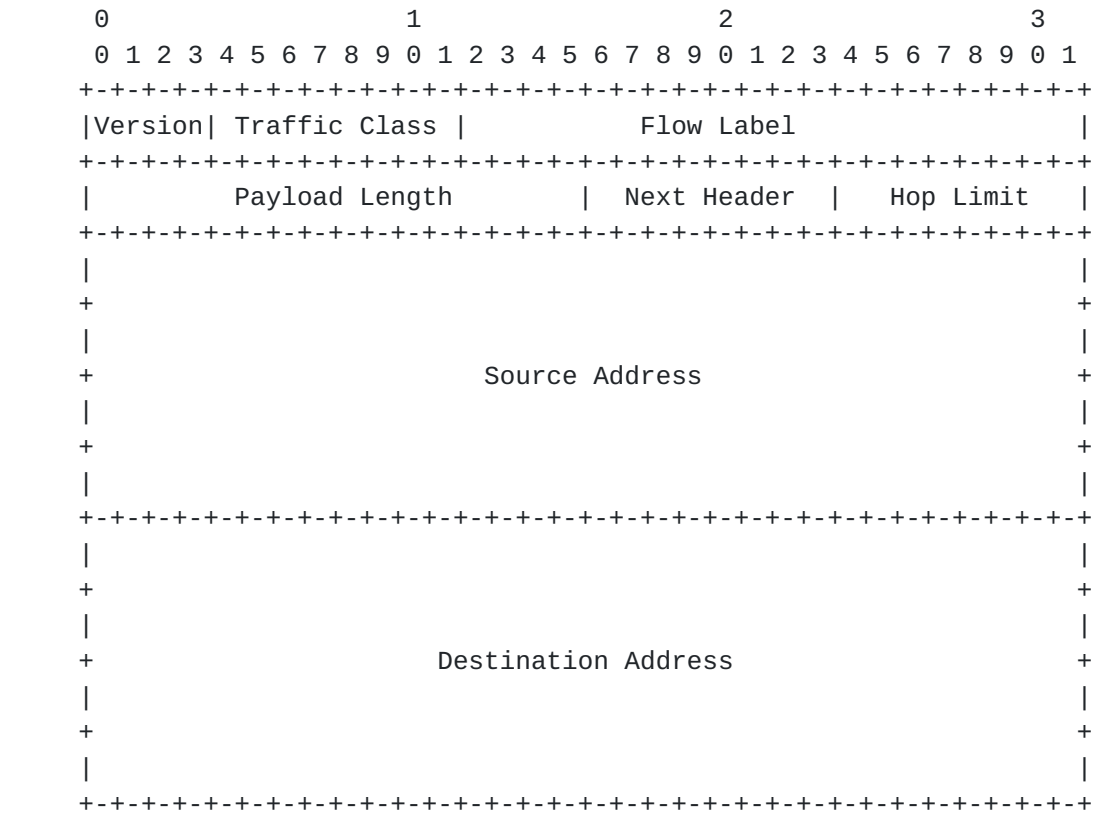

Figure 1: IPv6 Header

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 frequently viewed as imposing a burden on every IP router through which they pass.

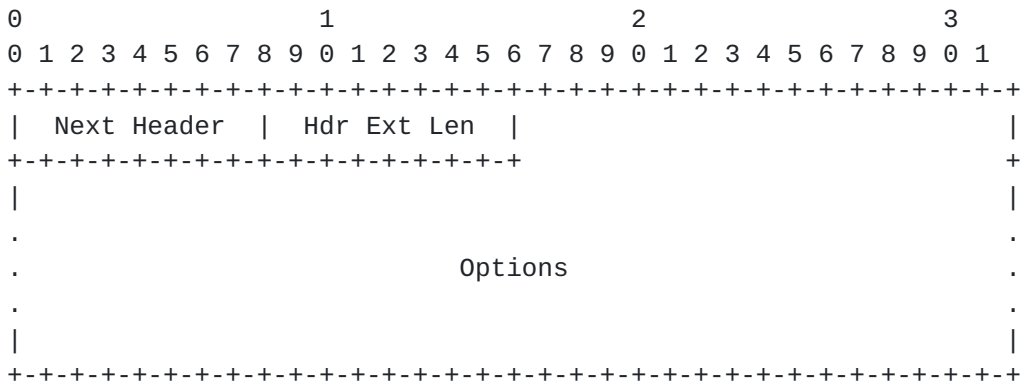

Figure 2: IPv6 Option Extension Header

2.2 IPv4 Options

Figure 3 shows the IPv4 header [RFC791]. The value of the initial 4-bit Version field indicates the IP version number and has value 4.

The IPv4 header has many similarities to the iPv6 header. For example, the IPv4 header 8-bit field called "Protocol" is like the "Next Header" field in the IPv6 header and the IPv4 header 8-bit "Type of Service" field, as amended by RFCs issued after [RFC791], is the same as the IPv6 header "Traffic Class" field. But options that are integrated into the more complex IPv4 header are handled by separate header extensions in IPv6. For example consider fragmentation, where an Internet Protocol packet is split into pieces, because the packet might be too big to traverse part of its

path, and these pieces are later recombined. Fragmentation is

indicated through an extension header for IPv6 but through fields in the main IPv4 header for IPv4. IPv4 options are considered part of the IPv4 header and the size of the options can be determined from the value of the IHL (Internet Header Length) field which gives the size of the IPv4 header in units of 4-octet words.

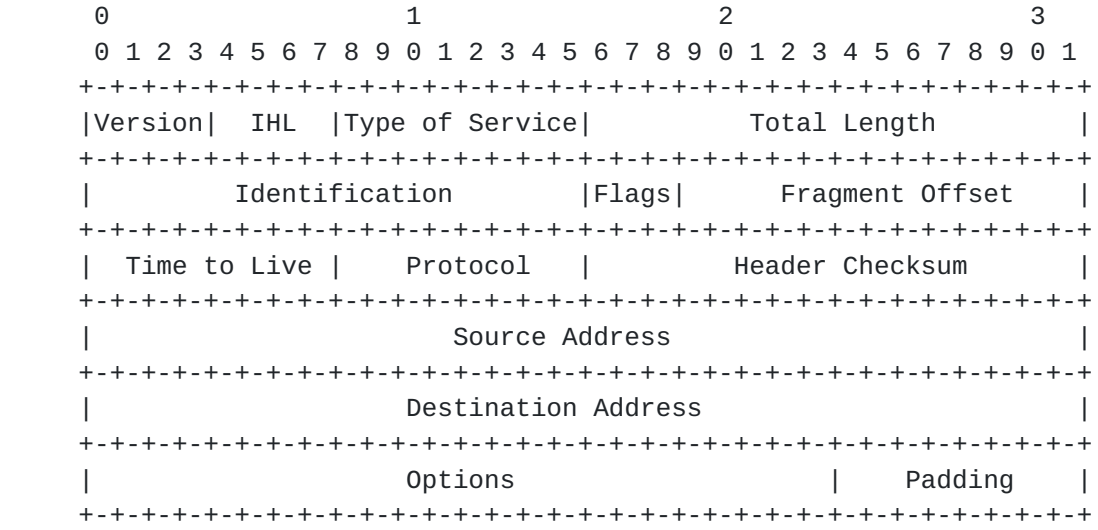

Figure 3: IPv4 Header

3. Overview of a Solution

Figure 4 shows a very high-level view of a network path between two hosts within local networks through the Internet core. (In reality there will 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 defalt free zone.)

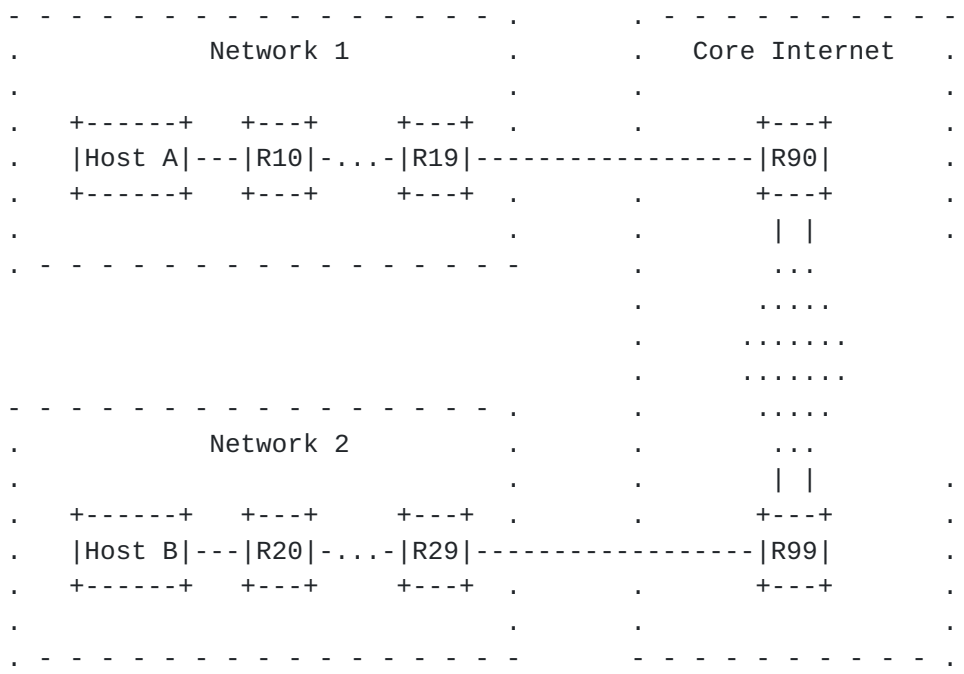

Figure 4: High Level View of an Internet Path

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 fully deployed in some network areas, there is likely to be substantial delay before it would be deployed in most of the Internet core. While some Internet core routers may ignore options, others discard all 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 hiding is done in an easily detectable and reversible fashion so that options can be unhidden after leaving the core. IPv6 Hop-by-Hop options or IPv4 options so hidden might not be effective in the core but the situation is 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 desireable 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. This transformation could also be applied to all packets starting with the first hsving 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. They are unhidden by changing this opaque IP protocol number in the IPv6 header back to zero. The points of hiding and unhiding 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 4) that implement the options seeing the options while minimizing he chance of unwanted packet discard.

The use of the opaque IP protocol number can defeat deeper IPv6 packet analysis that is intended to identify flows. It is therefore RECOMMENDED that, when this hiding technique is used, the IPv6 header Flow Label field be set [RFC6437] and used to identify flows [RFC6438] [RFC7098]. Using the Flow Label is a good idea anyway since IPv6 extension headers may move some fields on which flow identity might be based, such as port numbers, so deep into a packet that they are hard to use by routers.

3.2 Transiently Hiding IPv4 Options

A similar technique can be used for hiding IPv4 options but significantly more complex manipulations of the packet are required. As shown in Figure 5, the IPv4 header is made to appear to have no options by setting the IHL (Internet Header Length) field to its minimum value of 5, the Protocol field is changed to the opaque IP protocol number TBD, and the Header Checksum is adjusted to be correct for the optionless header. To be able to restore the IPv4 header, the old IHL, Protocol, and Header Checksum fields are saved in a 4-octet word inserted after the Destination Address and before any Options. The placement of the saved fields is such that their

alignment within a 4-octet word is the same as in the unmodified IPv4

header. The field labeled MBZ MUST be sent as zero and ignored on receipt.

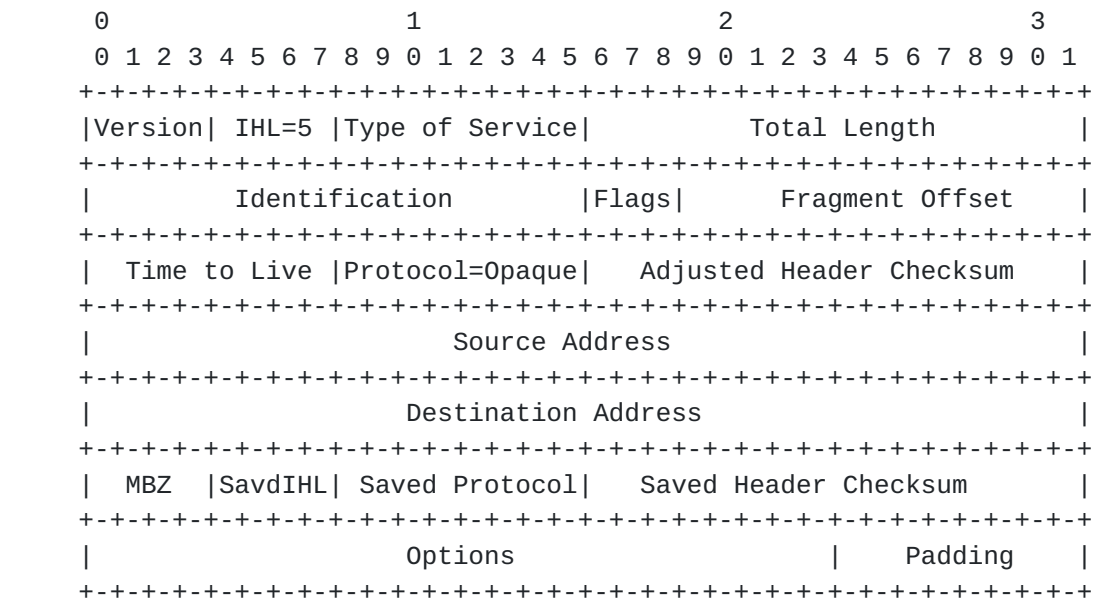

Figure 5: Modified IPv4 Header

These modifications increase the size of the IPv4 packet, increasing the chance that fragmentation or MTU problems could occur. For any node ignorant of the opaque IP protocol number, these modifications will interfere with flow determination based on the traditional 5-tuple (source and destination address, source and destination port, and IP protocol) or deep packet inspection.

3.3 Evolution to Greater Option Support

This solution supports the evolution of the Internet toward more widespread support of options as follows:

- As acceptable option support 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 support, 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 request 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]

5. Security Considerations

The use of the opaque IP Protocol to mask options is intended to defeat normal analysis of the following packet content, specifically options in the IP header. This would make firewalls, deep packet analysis, and the like less effective. On the other hand, firewalls tend to only admit packets with known permissable 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.

Should an IPv6 or IPv4 packet modified to hide options get through to a host, it would likely be discarded due to having an unknown IP Protocol.

More TBD

6. Acknowledgements

The helpful comments of the following are gratefully acknowledged:

Peng Shuping

Normative References

[RFC791] - Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, https://www.rfc-editor.org/info/rfc791

[RFC2119] - Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC6437] - Amante, S., Carpenter, B., Jiang, S., and J. Rajahalme, "IPv6 Flow Label Specification", RFC 6437, DOI 10.17487/RFC6437, November 2011, <https://www.rfc-editor.org/info/rfc6437>.

[RFC6438] - Carpenter, B. and S. Amante, "Using the IPv6 Flow Label for Equal Cost Multipath Routing and Link Aggregation in Tunnels", RFC 6438, DOI 10.17487/RFC6438, November 2011, <https://www.rfc-editor.org/info/rfc6438>.

[RFC8174] - Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <http://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

Informative References

[Options1] - Li, Z., Peng, S., and G. Mishra, "Hop-by-Hop Forwarding Options Header", Internet draft-li-6man-hbh-fwd-hdr-01, 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", Internet draft-hinden-6man-hbh-processing-01, June 2021, https://datatracker.ietf.org/doc/draft-hinden-6man-hbh-processing/

[Options3] - Peng, S., Li, Z., Xie, C., and Z. Qin, "Operational Issues with Processig of the Hop-by-Hop Options Header", Internet draft-ietf-v6ops-hbh-00, June 2021, https://datatracker.ietf.org/doc/draft-ietf-v6ops-hbh/

[RFC7098] - Carpenter, B., Jiang, S., and W. Tarreau, "Using the IPv6 Flow Label for Load Balancing in Server Farms", RFC 7098, DOI 10.17487/RFC7098, January 2014, <https://www.rfc-editor.org/info/rfc7098>.

Authors' Address

Donald E. Eastlake 3rd
Futurewei Technologies
2386 Panoramic Circle
Apopka, FL 32703 USA

Tel: +1-508-333-2270
Email: d3e3e3@gmail.com

Copyright and IPR Provisions