Workgroup: Network Working Group
Internet-Draft: draft-ietf-xml2rfc-template-06
Updates: 8200 (if approved)
Published: 7 July 2022
Intended Status: Standards Track
Expires: 8 January 2023
Authors: R. Hinden
Check Point Software
G. Fairhurst
University of Aberdeen

IPv6 Hop-by-Hop Options Processing Procedures

Abstract

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 of IPv6 Hop-by-Hop options practical with the goal of making IPv6 Hop-by-Hop options useful to deploy and use in the Internet. When published, this document updates RFC8200.

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 8 January 2023.

Table of Contents

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 of IPv6 Hop-by-Hop options practical with the goal of making IPv6 Hop-by-Hop options useful to deploy and use in the Internet.

The editors focus for this document is to set a lower bound expectation for the minimum number of hop-by-hop options that a node supports. This document does not discuss an upper bound. This topic is discussed in [I-D.ietf-6man-eh-limits].

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:

* Forwarding Plane: IPv6 hosts exchange user data through the forwarding plane. User data is processed by its recipient (i.e., an IPv6 host). User data can traverse intermediate nodes (i.e., routers) between its source and its destination. These intermediate nodes process metadata contained in packet headers. However, they do not process information contained in packet payloads.

* Control Plane: IPv6 routers exchange management and routing information with controllers. They also exchange routing information with one another. Management and routing information is processed by its recipient (i.e., an IPv6 router or controller). Management and control information can traverse intermediate nodes (i.e., routers) between its source and its destination. These intermediate nodes process metadata contained in packet headers. However, they do not process information contained in packet payloads. So, from their perspective, this information is user data.

* Fast Path: A path through a router that is optimized for forwarding packets without processing their payloads. The Fast Path may be supported by Application Specific Integrated Circuits (ASICS), Network Processor (NP), or other special purpose hardware. This is the usual processing path within a router taken by the forwarding plane.

* Slow Path: A path through a router that is capable of general purpose processing and is not optimized for any particular function. This is processing path is used for packets that require special processing or differ from assumptions made in Fast Path heuristics, or to process router control protocols used by the control plane.

NOTE: This distinct separation between hardware and software processing from [RFC6398] does not apply to all router architectures. However, a router that performs all or most processing in software might still incur more processing cost when providing special processing (aka Slow Path).

[RFC6192] is an example of how designs can separate control plane (Slow Path) and forwarding plane (Fast Path) functions.

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 full the forwarding rate (e.g., routers with no support on the Fast Path).

* Hop-by-Hop options would be sent to the Slow Path. This could could degrade the a router's performance and it's ability to process important control traffic.

* A mechanism that forces packets from any source to the routers "Slow Path" could be exploited as a Denial of Service attack against 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 were 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.

The text regarding processing Hop-by Hop Options in [RFC8200] was not intended to change the processing of Hop-by-Hop options. It only documented how they were being used in the Internet at the time

RFC8200 was published. This was a constraint on publishing the IPv6 specification as an IETF Standard.

The main issues remain:

* Routers are commonly configured to drop transit packets containing hop-by-hop options that would have be processed in the Slow Path. This behavior is seen as protecting against a denial of service attack on the router. A survey in 2015 reported a high loss rate in transit ASs for packets with HBH options [RFC7872]. The operational implications of IPv6 Packets that set extension headers is discussed in [RFC9098].

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

* Any mechanism that can be used 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 needed 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 that 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 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 Extenstion Headers we therefore recommend to support at least one EH".

The authors expectations are that some hop-by-hop options will be processed across the Internet while others will only be processed in a limited domain (e.g., where there is a specific service made

available in that network segment that relies on one or more hop-by-hop options).

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 several 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 of [RFC8200] 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] that a node MUST process the first Option in the Hop-by-Hop Header at full forwarding rate the (e.g. on the router's Fast Path) and MAY process additional Hop-by-Hop Options if configured to do so. The motivation for this change is to simplify the processing of Hop-by-Hop options as a part of normal forwarding.

Nodes creating packets with a Hop-by-Hop option headers SHOULD include a single Hop-by-Hop Option in the packet and MAY include more based on local configuration.

If there are more than one Hop-by-Hop options in the Hop-by-Hop Options header, the node MAY skip the rest of the options without having to examine these options using the "Hdr Ext Len" field in the Hop-by-Hop Options header. This field specifies the length of the Option Header in 8-octet units. The additional options do not need to be processed or verified.

5.2. Hop-by-Hop Headers Processing

Nodes that implement a differentiation between a Fast Path and a Slow Path MUST process all (with one exception noted below) Hop-by-Hop options in the Fast Path. The one exception to this is the Router Alert Option [RFC2711]. See Section 5.3 for discussion of the Router Alert.

If the node can not process an option at the full forwarding rate, it MUST behave as if it does not recognize the Option Type (as described in the next paragraph).

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. The text is:

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

01 - discard the packet.

10 - discard the packet and, regardless of whether or not the packet's Destination Address was a multicast address, send an ICMP Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.

11 - discard the packet and, only if the packet's Destination Address was not a multicast address, send an ICMP Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.

This document modifies this behaviour for the "10" and "11" values that the node MAY send an ICMP Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type. The modified text for "10" and 11" values is:

10 - discard the packet and, regardless of whether or not the packet's Destination Address was a multicast address, MAY send an ICMP Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.

11 - discard the packet and, only if the packet's Destination Address was not a multicast address, MAY send an ICMP Parameter Problem, Code 2, message to the packet's Source Address, pointing to the unrecognized Option Type.

The motivation for this change is to loosen the requirement to send ICMPv6 Parameter Problem messages by simplifying what the router needs to do when it performs forwarding of an Option Type it does not recognize.

When an ICMP Parameter Problem, Code 2, message is delivered to the source, the source can become aware that at least one node on the path has failed to recognize the option.

5.3. Router Alert Option

The Router Alert option [RFC2711] purpose is to tell the node that the packet needs additional processing on the Slow Path.

The Router Alert option includes a two octet Value field that describes the protocol that is carried in the packet. The current values can be found in the IANA Router Alert Value registry [IANA-RA].

DISCUSSION

> The Router Alert Option is a problem since it's function is to do what this specification is proposing to eliminate, that is, process the packet in the Slow Path. One approach would be to deprecate it as it's usage appears to be limited and packets containing Hop-by-Hop options are frequently dropped. Deprecation would allow current implementations to continue and it's use could be phased out over time.
>
> The authors current thinking is that the Router Alert function may have reasonable potential use for new functions that have to be processed in the Slow Path. We think that keeping it as the single exception for Slow Path processing with the following restrictions is a reasonable compromise to allow future flexibility. These are compatible with Section 5 of [RFC6398].

A Fast Path implementation SHOULD verify that a Router Alert contains a protocol, as indicated by the Value field in the Router Alert option, that is configured as a protocol of interest to that router. A verified packet SHOULD be sent on the Slow Path for processing [RFC6398]. Otherwise, the router implementation SHOULD forward within the Fast Path (subject to all normal policies and forwarding rules). As specified in [RFC2711] the top two bits of Option Type for the Router Alert option are always set to "00" indicating the node should skip over this option and continue processing the header in this case.

Implementations of the IP Router Alert Option SHOULD offer the configuration option to simply ignore the presence of "IP Router Alert" in IPv4 and IPv6 packets" [RFC6398].

A node that is configured to process a Router Alert option using the Slow Path MUST protect itself from infrastructure attack that could result from processing on the Slow Path. This might include some combination of access control list to only permit from trusted nodes, rate limiting of processing, or other methods [RFC6398].

5.4. 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.

A node configured not to process HBH options, MUST drop the packet if the top two bits of the Option Type field of the first HBH option is non-zero.

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

6. New Hop-by-Hop Options

Any new IPv6 Hop-by-Hop option designed in the future should be designed to be processed at full forwarding rate (e.g., on a router's Fast Path). New options SHOULD NOT be defined that are not expected to be executed at full forwarding rates. New Hop-by-Hop options SHOULD have the following characteristics:

* Straight forward to process. That is, they should be designed to keep the time to process low.

* New Hop-by-Hop options should be designed to be the first option in a Hop-by-Hop options header.

* The size of an option should not extend beyond what can be reasonably expected to be executed at full forwarding rate (e.g., forwarded on a router's fast path).

Any new Hop-by-Hop option that is standardized that does not meet these criteria needs to explain in detail in its specification why this can not be accomplished and that there is a reasonable expectation that it can be proceed at full forwarding rate.

7. IANA Considerations

There are no actions required for IANA defined in this document.

8. Security Considerations

Security issues with IPv6 Hop-by-Hop options are well known and have been documented in several places, including [RFC6398], [RFC6192], and [RFC9098]. The main issue, as noted in Section 4, is that any mechanism that can be used 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. Due to this it's common for transit routers to drop packets with Hop-by-Hop options headers.

While Hop-by-Hop options are not required to be processed in the Slow Path, the Router Alert options is designed to do just that.

This document changes the way Hop-by-Hop options are processed in several ways that significantly reduces the attack surface. These changes include:

* All Hop-by-Hop options (with one exception) must be processed in the Fast Path. Only one HBH Option MUST be processed and additional HBH Options MAY be processed based on local configuration.

* Only the Router Alert option can be processed in the Slow Path, and the router must be configured to do so.

* Added criteria to allow control over how Router Alert options are processed and that a node configured to support these options must protect itself from attacks using the Router Alert.

* Limited the default number of Hop-by-Hop options that that can be in a packet to a single Hop-by-Hop option.

* Additional Hop-by-Hop options MAY be included, based on local configuration. Although nodes only process these additional Hop-by-Hop Options if configured to do so.

* Added restrictions to any future new Hop-by-Hop options that limit their size and computational requirements.

The authors believe that these changes significantly reduces the security issues relating to IPv6 Hop-by-Hop options and will enable them to be used safely in the Internet.

9. Acknowledgments

Helpful comments were received from Brian Carpenter, Ron Bonica, Ole Troan, Mark Heard, Tom Herbert, Cheng Li, Eric Vyncke, Greg Mirksy, Xiao Min, Fernando Gont, Darren Dukes, [your name here], and other members of the 6MAN working group.

10. Change log [RFC Editor: Please remove]

draft-ietf-6man-hbh-processing-01, 2022-June-15:

- *Fixed typo in last paragraph of Section 5.2
- *Revised text in Section 4 to reflect constraints on publishing RFC8200.
- *Changed text in Section 6 that new options SHOULD NOT (from MUST NOT) be defined that require that are not expected to be excepted at full forwarding rates.
- *Added reference to RFC7872 in Section 4.
- *Added text to Section 1 that the focus of this document is to set a minium bound on the number of Hop-by-Hop Options a node should process.
- *Added text to Section 4 that the authors some Hop-by-Hop optoions will be supported Internet wide, and others only in limited domains.
- *Editorial changes.

draft-ietf-6man-hbh-processing-00, 2022-January-29:

- *6MAN Working Group Draft
- *Reworked text to talk about processing HBH options at full forwarding rates, instead of "fast path"
- *Revised Section 6 "New Hop-by-Hop Options" to allow variable sized HBH options, remove specific length requirements, and other clarifications.
- *Editorial changes.

draft-hinden-6man-hbh-processing-01, 2021-June-2:

- *Expanded terminology section to include Forwarding Plane and Control Plane.
- *Changed draft that only one HBH Option MUST be processed and additional HBH Options MAY be processed based on local configuration.
- *Clarified that all HBH options (with one exception) must be processed on the Fast Path.
- *Kept the Router Alert options as the single exception for Slow Path processing.
- *Rewrote and expanded section on New Hop-by-Hop Options.
- *Removed requirement for HBH Option size and alignment.
- *Removed sections evaluating currently defined HBH Options.

*Added content to the Security Considerations section.
*Added people to the acknowledgements section.
*Numerous editorial changes

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

*Initial draft.

11. 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>.

12. Informative References

[Hendriks] Hendriks, L., Velan, P., Schmidt, RO., 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>.

[I-D.ietf-6man-eh-limits]
Herbert, T., "Limits on Sending and Processing IPv6 Extension Headers", Work in Progress, Internet-Draft, draft-ietf-6man-eh-limits-00, 31 January 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-6man-eh-limits-00>.

[IANA-RA] "IPv6 Router Alert Option Values", <https://www.iana.org/assignments/ipv6-routeralert-values/ipv6-routeralert-values>.

[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>.

[RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, DOI 10.17487/RFC6192, March 2011, <https://www.rfc-editor.org/info/rfc6192>.

[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>.

[RFC7872] Gont, F., Linkova, J., Chown, T., and W. Liu, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", RFC 7872, DOI 10.17487/RFC7872, June 2016, <https://www.rfc-editor.org/info/rfc7872>.

[RFC9098] Gont, F., Hilliard, N., Doering, G., Kumari, W., Huston, G., and W. Liu, "Operational Implications of IPv6 Packets with Extension Headers", RFC 9098, DOI 10.17487/RFC9098, September 2021, <https://www.rfc-editor.org/info/rfc9098>.

Authors' Addresses

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

Email: 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