IPv6 Maintenance
Internet-Draft
Updates: 2460,7045 (if approved)
Intended status: Standards Track
Expires: September 17, 2016

F. Baker
Cisco Systems
R. Bonica
Juniper Networks

March 16, 2016

IPv6 Hop-by-Hop Options Extension Header
draft-ietf-6man-hbh-header-handling-03

Abstract

This document clarifies requirements for IPv6 routers with respect to the Hop-by-Hop (HBH) Options Extension Header. These requirements are applicable to all IPv6 routers, regardless of whether they maintain a strict separation between forwarding and control plane hardware. In this respect, this document updates RFC 2460 and RFC 7045.

This document also describes forwarding plane procedures for processing the HBH Options Extension Header. These procedures are applicable to implementations that maintain a strict separation between forwarding and control plane implementations.

The procedures described herein satisfy the above mentioned requirements by processing HBH Options on the forwarding plane to the greatest degree possible. If a packet containing HBH Options must be dispatched to the control plane, it is rate limited before dispatching. In order to comply with the requirements of this specification, implementations may execute the procedures described herein or any other procedures that result in compliant behavior.

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 http://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 September 17, 2016.

Table of Contents

1. Introduction

In IPv6 [RFC2460], optional Internet-layer information is encoded in extension headers that may be placed between the IPv6 header and the upper-layer header. Currently, eleven extension headers are defined. Among them is the Hop-by-Hop (HBH) Options Extension header. Unlike any other extension header, the HBH Options Extension header is examined by every node that a packet visits en route to its destination.

The HBH Extension Header contains one or more HBH Options. Each HBH Option contains a type identifier. Appendix B of this document provides a list of currently defined HBH options.

Some HBH Options contain information that is useful to a router's forwarding plane. In this document, we call these options "HBH forwarding options". Among these is the Jumbo Payload Option

[RFC2675]. The Jumbo Payload Option indicates the payload length of the packet that carries it. While this information is required to forward the packet, it can be discarded as soon as the packet has been forwarded.

By contrast, other HBH Options contain information that is useful to a router's control plane. In this document, we call these options "HBH control options". Among these is the Router Alert Option [RFC2711]. The Router Alert Option informs transit routers that the packet carrying it contains information to be consumed by the router's control plane. In many cases, this information is used to forward subsequent packets.

Finally, the Pad and Pad1 options contain no information at all. These are included to ensure word-alignment of subsequent options and headers.

Many modern routers maintain a strict separation between forwarding plane hardware and control plane hardware. In these routers, forwarding plane bandwidth is plentiful, while control plane bandwidth is constrained. In order to protect scarce control plane resources, these routers enforce policies the restrict access from the from the forwarding plane to the control plane. Effective policies address packets containing the HBH Options Extension header, because HBH control options require access from the forwarding plane to the control plane.

Many network operators perceive HBH Options to be a breach of the separation between the forwarding and control planes [I-D.ietf-v6ops-ipv6-ehs-in-real-world]. Therefore, some network operators discard all packets containing the HBH Options Extension Header, while others forward the packets but ignore the HBH Options. Still other operators severely rate-limit packets containing the HBH Options Extension Header. In addition, some (notably older) implementations send all packets containing a HBH header to the control plane even if they contain only pad options, resulting in an effect DoS on the router and inconsistent drops among those packets due to rate limiting or other factors.

[RFC7045] legitimizes the current state of affairs, severely limiting the utility of HBH options. In the words of RFC 7045:

> "The IPv6 Hop-by-Hop Options header SHOULD be processed by intermediate forwarding nodes as described in RFC2460. However, it is to be expected that high-performance routers will either ignore it or assign packets containing it to a slow processing path. Designers planning to use a Hop-by-Hop option need to be aware of this likely behaviour."

This document clarifies requirements for IPv6 routers with respect to the HBH Options Extension Header. These requirements are applicable to all IPv6 routers, regardless of whether they maintain a strict separation between forwarding and control plane hardware. In this respect, this document updates RFC 2460 and RFC 7045.

This document also describes forwarding plane procedures for processing the HBH Options Extension Header. These procedures are applicable to implementations that maintain a strict separation between forwarding and control plane hardware.

The procedures described herein satisfy the above mentioned requirements by processing HBH Options on the forwarding plane to the greatest degree possible. If a packet containing HBH Options must be dispatched to the control plane, it is rate limited before dispatching. In order to comply with the requirements of this specification, implementations can execute the procedures described herein or any other procedures that result in compliant behavior.

1.1. 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 [RFC2119].

2. Requirements

This section clarifies requirements for IPv6 routers with respect to the HBH Options Extension Header. These requirements are applicable to all IPv6 routers, regardless of whether they maintain a strict separation between forwarding and control plane hardware.

- REQ1: Implementations MUST NOT discard otherwise forwardable packets because they contain the HBH Options Extension header. However, an implementation MAY be configured to discard packets containing the HBH Options Extension Header, so long as this is not the default behavior.

- REQ 2: Implementations MUST process unrecognized HBH Options as described in Section 4.2 of RFC 2460. If an implementation receives a packet that contains an unrecognized HBH Option, that implementation MUST examine the first two bits of the HBH Option Type indicator. Those bits determine whether the implementation a) continues to process the packet, b) discards the packet without sending an ICMP message or c) discards the packet and sends an ICMP message.

- REQ 3: Unrecognized HBH Options MUST be evaluated sequentially. For example, assume that an implementation receives a packet that carries two unrecognized HBH Options. The Type indicator of the first unrecognized option begins with 01 while the Type indicator of the second unrecognized option begins with 10. In this case, the implementation MUST discard the packet without sending an ICMP message to the source. However, if the Type indicator of the first unrecognized option begins with 10 and the Type indicator of the second unrecognized option begins with 01, the implementation MUST discard the packet and send and ICMP Parameter Problem message to the source.

- REQ 4: Implementations MUST protect themselves against denial of service attacks that are propagated through HBH Options. These protections MUST be enabled by default, without special configuration.

- REQ 5: The originator of a packet MAY insert the HBH Options Extension header between the IPv6 header and the upper-layer header. It MAY also insert HBH Options inside of the HBH Options header. Transit routers MUST NOT insert the HBH Options Extension header between the IPv6 header and the upper-layer header. Furthermore, they MUST NOT add or delete HBH Options inside of the HBH Options Extension header.

- REQ 6: Implementations SHOULD support a configuration option that limits the set of HBH Options that they recognize. For example, assume that an implementation recognizes a particular HBH Option. Using this configuration option, an operator can cause the implementation to behave as if it does not recognize that option. This MAY be configured a side effect of other functionality. For example, an implementation might not recognize the Router Alert Option unless a protocol that relies on the Router Alert Option (e.g., RSVP) is configured.

- REQ 7: The HBH Options Extension Header can contain as many as 2056 bytes. Some implementation are not capable of processing extension headers of that length [I-D.gont-v6ops-ipv6-ehs-packet-drops]. When an implementation receives a packet that it cannot process due to its HBH Options Extension Header length, the implementation MUST discard the packet and send an ICMP Parameter Problem message to the packet source. ICMP Parameter Problem Code MUST be "Long Extension Header" (value TBD) and the ICMP Parameter Problem Pointer MUST contain the offset of HBH Options Extension Header.

3. Proposed Procedures

This section describes forwarding plane procedures for processing the HBH Options Extension Header. These procedures are applicable to implementations that maintain a strict separation between forwarding and control plane hardware.

The procedures described below process HBH Options on the forwarding plane to the greatest degree possible. If a packet containing HBH Options must be dispatched to the control plane, it is rate limited before dispatching. In order to comply with the requirements of Section 2, implementations can execute the procedures described herein or any other procedures that result in compliant behavior.

Having received a packet containing the HBH Options Extension header, the forwarding plane determines whether the HBH Options Extension Header is too long for it to process. If so, the forwarding plane discards the packet and sends an ICMP Parameter Problem message to the packet source. ICMP Parameter Problem Code is set to "Long Extension Header" and the ICMP Parameter Problem Pointer is set to the offset of HBH Options Extension Header.

If the HBH Options Extension Header is not too long to process, the forwarding plane hardware scans the header, assigning it to one of the following classes:

- Discard
- Dispatch to control plane
- Forward, ignoring all HBH Option
- Forward, processing selected HBH Options

Forwarding plane hardware discards the packet if the HBH Options Extension Header contains an unrecognized option whose Type indicator begins with 01, 10 or 11. Forwarding plane hardware sends an ICMP message if required. See Section 2 REQ 2 and REQ 3 for details.

If the packet is not discarded, and the HBH Options Extension header contains at least one recognized control option, the forwarding plane subjects the packet to a rate-limit and dispatches it to the control plane

Otherwise, if the HBH Options Extension header contains only the following option types, the packet is forwarded without further HBH Option processing:

- Pad or Pad1
- Unrecognized options whose Type indicator begins with 00

Otherwise, the forwarding plane process forwarding options and forwards the packet

4. IANA Considerations

IANA is requested to assign a new entry to the ICMP Parameter Problem Code registry. The name of this code is "Long Extension Header".

5. Security Considerations

This document contributes to the security of IPv6 routers, by defining forwarding plane procedures for the processing of HBH Options. These procedures are applicable to implementations that maintain a strict separation between forwarding and control plane hardware.

The procedures described below process HBH Options on the forwarding plane to the greatest degree possible. If a packet containing HBH Options must be dispatched to the control plane, it is rate limited before dispatching.

6. Acknowledgements

This note grew out of a discussion among the author, Ole Troan, Mark Townsley, Frank Brockners, and Shwetha Bhandari, and benefited from comments by Dennis Ferguson, Brian Carpenter, Panos Kampanakis, Jinmei Tatuya, and Joe Touch. Thanks to Fernando Gont for his thoughtful review.

7. References

7.1. 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, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <http://www.rfc-editor.org/info/rfc2460>.

[RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", RFC 7045, DOI 10.17487/RFC7045, December 2013, <http://www.rfc-editor.org/info/rfc7045>.

7.2. Informative References

[I-D.gont-v6ops-ipv6-ehs-packet-drops]
Gont, F., Hilliard, N., Doering, G., LIU, S., and W. Kumari, "Operational Implications of IPv6 Packets with Extension Headers", draft-gont-v6ops-ipv6-ehs-packet-drops-03 (work in progress), March 2016.

[I-D.ietf-6man-rfc2460bis]
Deering, S. and B. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", draft-ietf-6man-rfc2460bis-03 (work in progress), January 2016.

[I-D.ietf-roll-trickle-mcast]
Hui, J. and R. Kelsey, "Multicast Protocol for Low power and Lossy Networks (MPL)", draft-ietf-roll-trickle-mcast-12 (work in progress), June 2015.

[I-D.ietf-v6ops-ipv6-ehs-in-real-world]
Gont, F., Linkova, J., Chown, T., and S. LIU, "Observations on the Dropping of Packets with IPv6 Extension Headers in the Real World", draft-ietf-v6ops-ipv6-ehs-in-real-world-02 (work in progress), December 2015.

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

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

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, DOI 10.17487/RFC4302, December 2005, <http://www.rfc-editor.org/info/rfc4302>.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, DOI 10.17487/RFC4303, December 2005, <http://www.rfc-editor.org/info/rfc4303>.

[RFC4782] Floyd, S., Allman, M., Jain, A., and P. Sarolahti, "Quick-Start for TCP and IP", RFC 4782, DOI 10.17487/RFC4782, January 2007, <http://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, <http://www.rfc-editor.org/info/rfc5570>.

[RFC6398] Le Faucheur, F., Ed., "IP Router Alert Considerations and Usage", BCP 168, RFC 6398, DOI 10.17487/RFC6398, October 2011, <http://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, <http://www.rfc-editor.org/info/rfc6553>.

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

[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, <http://www.rfc-editor.org/info/rfc6971>.

Appendix A. Change Log

RFC Editor: this section need not be published in any RFC.

Initial Version: October 2015: text copied from draft-baker-6man-hbh-header-handling-03.txt and discussed in IETF 94

IETF 94 Update: Sections 2.2, 2..3, and 2.4 moved to an appendix reflecting (negative) working group viewpoint on the modification of packet length in flight.

The content of this document is likely to be subsumed into 2460bis [I-D.ietf-6man-rfc2460bis], but is held separate for the present discussion.

A new section 2.2 added detailing conceptual processing model for HBH options.

version 2 Addressed editorial comments

Appendix B. HBH Options

At this writing, there are several defined Hop-by-Hop options:

PAD Options: The PAD1 and PADn [RFC2460]

Router Alert Option: The IPv6 Router Alert Option [RFC2711] [RFC6398]

Jumbo Payload: [RFC2675]

RPL Option: [RFC6553]

Quickstart Option [RFC4782]

Common Architecture Label IPv6 Security Option: [RFC5570]

SMF Option: [RFC6621]

MPL Option: [I-D.ietf-roll-trickle-mcast]

DFF Option: [RFC6971]

Authors' Addresses

Fred Baker
Cisco Systems
Santa Barbara, California 93117
USA

Email: fred@cisco.com

Ron Bonica
Juniper Networks
Herndon, Virginia 20171
USA

Email: rbonica@juniper.net