

Workgroup: 6man
Internet-Draft:
draft-bonica-6man-ext-hdr-update-07
Updates: [RFC 8200](#) (if approved)
Published: 24 February 2022
Intended Status: Standards Track
Expires: 28 August 2022
Authors: R. Bonica T. Jinmei
 Juniper Networks Infoblox

Inserting, Processing And Deleting IPv6 Extension Headers

Abstract

This document provides guidance regarding the processing, insertion, and deletion of IPv6 extension headers. It updates RFC 8200.

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 28 August 2022.

Copyright Notice

Copyright (c) 2022 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. Terminology](#)
- [3. Updates To RFC 8200](#)
 - [3.1. Original Text](#)
 - [3.2. Updated Text](#)
- [4. Motivation](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
- [7. Acknowledgements](#)
- [8. Normative References](#)
- [Authors' Addresses](#)

1. Introduction

In [IPv6](#) [[RFC8200](#)] optional internet-layer information is encoded in extension headers. As specified by [[RFC8200](#)], "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 statement quoted above identifies nodes upon which extension headers are not processed, inserted, or deleted. It does not imply that extension headers can be processed, inserted, or deleted on any other node along a packet's delivery path.

This document provides guidance regarding the processing, insertion, and deletion of IPv6 extension headers. It clarifies the statement quoted above and updates [[RFC8200](#)].

2. Terminology

The following terms are used in this document:

*Source node - An IPv6 source node accepts data from an upper-layer protocol, prepends an IPv6 header, and sends the resulting IPv6 packet to a destination node.

*Final destination node - An IPv6 final destination node receives an IPv6 packet and delivers its payload to an upper-layer protocol.

*Delivery path - A packet's delivery path is a series of nodes that a packet traverses on route to its final destination. The delivery path includes the final destination node.

*Segment - A segment is a series of links and nodes in a packet's delivery path. An IPv6 Routing header steers packets from segment to segment along the delivery path. If a packet contains a Routing header, its delivery path can contain multiple segments. If a packet does not contain a Routing header, its delivery path contains only one segment.

*Segment egress node - A segment egress node terminates a segment. When a packet arrives at a segment egress node, its IPv6 Destination Address identifies an interface that belongs to the node. All final destination nodes are also segment egress nodes.

*Extension header processing - Each IPv6 extension header is associated with a procedure. For example, the Fragment header is associated with fragmentation and reassembly procedures. Extension header processing is the reception of an extension header and the execution of its associated procedure.

3. Updates To RFC 8200

The terms defined in [Section 2](#) of this document should be added to Section 2 of [[RFC8200](#)].

[Section 3.1](#) of this document quotes text from [[RFC8200](#)]. That text should be replaced with the text contained by [Section 3.2](#) of this document.

3.1. Original Text

"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."

3.2. Updated Text

Source nodes can send packets that include extension headers. Extension headers are not inserted by subsequent nodes along a packet's delivery path.

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.

The Hop-by-Hop Options header can be processed by any node in a packet's delivery path. All remaining extension headers can be processed at segment egress nodes only. While some extension headers are processed at any segment egress node, others (e.g., the Fragment header) can only be processed at the final destination node.

Except for the Routing header, extension headers cannot be deleted by any node along a packet's delivery path. If the following conditions are true, a Routing header can be deleted by any segment egress node:

- *The Segments Left field in the routing header is equal to zero.

- *The packet does not contain an Authentication header.

Extension headers can be inspected for various purposes (e.g., firewall filtering) by any node along a packet's delivery path.

4. Motivation

The following are reasons why extension headers are not inserted by nodes along a packet's delivery path:

- *Nodes that execute [Path MTU Discovery \(PMTUD\)](#) [[RFC8201](#)] procedures can send packets that are nearly as large as the Path MTU. Adding an extension header to such a packet can cause MTU black holing.

- *[IPv6 Authentication Header](#) [[RFC4302](#)] processing relies on the immutability of the Payload Length field in the IPv6 header. When a node along a packet's delivery path inserts an extension header, it must also update the Payload Length field in the IPv6 header. Therefore, it causes IPv6 Authentication Header processing to fail on the final destination node.

- *When a source node sends a packet to a final destination node, and a node along the packet's delivery path inserts an extension header, the final destination node will mistakenly attribute the extension header to the source node. Attackers can leverage this mistaken attribution.

The following are reasons why extension headers, except for the Routing header, are not deleted by any node along a packet's delivery path:

*IPv6 Authentication Header processing relies on the immutability of the Payload Length field in the IPv6 header. When a node along a packet's delivery path inserts an extension header, it must also update the Payload Length field in the IPv6 header. Therefore, it causes IPv6 Authentication Header processing to fail on the final destination node.

*When a source node sends a packet to a final destination node, and a node along the packet's delivery path removes an extension header, the resulting packet may not elicit the behavior intended by the source node. For example, if a Destination Options header is removed, none of the options that it contains will be delivered to the final destination node.

The following are reasons why Routing headers can be deleted by any segment egress node when the Segments Left field is equal to zero and the packet does not contain an authentication header:

*Because every segment that the routing header contains has already been processed.

*Because [[RFC8986](#)] has set a precedent for deletion in this case.

5. Security Considerations

This document does not introduce any new security considerations.

6. IANA Considerations

This document does not request any IANA actions.

7. Acknowledgements

Thanks to Bob Hinden, Brian Carpenter, Tom Herbert and Fernando Gont for their comments and review.

8. Normative References

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

[[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>>.

[RFC8201]

McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed.,
"Path MTU Discovery for IP version 6", STD 87, RFC 8201,
DOI 10.17487/RFC8201, July 2017, <<https://www.rfc-editor.org/info/rfc8201>>.

[RFC8986]

Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer,
D., Matsushima, S., and Z. Li, "Segment Routing over IPv6
(SRv6) Network Programming", RFC 8986, DOI 10.17487/
RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.

Authors' Addresses

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, Virginia 20171
United States of America

Email: rbonica@juniper.net

Tatuya Jinmei
Infoblox

Email: jinmei@wide.ad.jp