Network Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: 11 July 2026

T. Herbert
XDPnet

7 January 2026

Deprecating IPv6 Extension Headers on the Internet
draft-herbert-deprecate-eh-01

Abstract

This document describes the deprecation of IPv6 extension headers on the Internet with the exception of Encapsulating Security Payload. Deprecation is motivated by three factors: 1) the data shows high discard rates for packets with extension headers sent over the Internet, 2) extension headers can be used for Denial of Service attack and are replete with other security vulnerabilities, 3) the high loss rates are a disincentive to develop new extension headers or options that might be useful or fix known problems. This document recommends that extension headers, other than Encapsulating Security Payload, be relegated to use only in limited domains and that packets with extension headers should be discarded at boundary routers of limited domains.

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 11 July 2026.

Copyright Notice

Table of Contents

1. Introduction 2
1.1. Problem statement 3
1.2. Requirements Language 4
2. Requirements for different extension headers 4
2.1. Hop-by-Hop Options 4
2.2. Routing Header 5
2.3. Encapsulating Security Payload (ESP) 5
2.4. Destination Options 6
2.5. Fragment Header 6
2.6. Authentication Header 7
2.7. Sending ICMP errors 7
3. Requirements for extension headers in limited domains 8
4. Acknowledgments 8
5. IANA Considerations 8
6. Security Considerations 8
7. References 8
7.1. Normative References 8
7.2. Informative References 8
Author's Address 10

1. Introduction

Extension headers were ostensibly defined to make IPv6 extensible. However, after twenty-five years plus of history with IPv6, extension headers have not been widely deployed and cannot be reliably sent over the Internet. Arguably, extension headers are a failed experiment in protocol design. The open-ended nature of extension headers has a lot to do with that, but there are also fundamental problems in security considering that extension headers are sent in plain text with no integrity, confidentiality, or even a simple checksum to guard against data corruption like the IPv4 header checksum protects IP Options.

A fix to the extension Header quagmire, at least the part caused by defining extension headers as an open ended protocol, was proposed in [eh-limits]. Unfortunately, that draft was rejected because of a belief that it would ossify the protocol and some people didn't like

the term "limits". Given that no one else has proposed a fix for the problems of extension headers and no other alternatives seem to be forthcoming, we believe that deprecating them, at least for use on the Internet, is the most prudent course of action.

This document proposes that extension headers be deprecated on the Internet by discarding packets with extension headers at the boundaries of limited domains [RFC8799]. Given the high loss rates of extension headers on the Internet and the fact that using extension headers is considered a non-starter for application developers, this specification is really just codifying reality.

1.1. Problem statement

Extension headers are a core component of the IPv6 protocol as specified in [RFC8200]. IPv6 extension headers were originally defined with few restrictions. For instance, there is no specified limit on the number of extension headers a packet may have, nor is there a limit on the length in bytes of extension headers in a packet other than being limited by the Path MTU or 1,280 bytes for those hosts that do not discover the Path MTU [RFC7112]. Similarly, variable length extension headers typically do not have prescribed limits such as limits on the number of Hop-by-Hop or Destination options in a packet. The lack of limits essentially requires implementations to handle every conceivable usage of the protocol, including myriad use cases outside the realm of ever being realistic or useful in real world deployment.

The lack of limits and the requirements for supporting a virtually open-ended protocol have led to a current lack of support and deployment of extension headers ([RFC7872], [Cus23b]). Instead of attempting to satisfy the protocol requirements concerning extension headers, some router and middlebox vendors have opted to invent and apply their own ad hoc limits, relegate packets with extension headers to slow path processing, or have gone so far as to summarily discard all packets with extension headers [RFC9098]. For those hosts and routers that properly attempt to process all extension headers per the specifications, the lack of limits has made them susceptible to Denial of Service attacks. The net effect of this situation is that deployment and use of extension headers is currently underwhelming. [Cus23a] and [APNIC] provide data on the drop rates of extension headers on the Internet.

In addition to the lack of limits for extension headers, concerns about extension headers have been raised with regard to their susceptibility to Denial of Service attack and security vulnerabilities. Except for the case where extension headers are encrypted, sending any information in or accepting any information in

plain text on an untrusted path that is not necessary for routing packets is an obvious security vulnerability that risks information leakage or spoofing of the information. Extension headers also facilitate various Denial of Service attacks especially with Hop-by-Hop Options and Routing Header since those target network infrastructure and not just end hosts. The security issues with IPv6 extension headers are discussed in [RFC4942].

The high loss rates and other issues of extension headers are a demotivation for would-be protocol developers to develop new extension headers or options. For instance, after more than twenty-five years of IPv6 there are no Destination or Hop-by-Hop options that one would be considered universally useful, much less required. This creates a canonical "chicken and the egg problem": network administrators won't enable extension headers in their network without seeing evidence of useful options; developers won't develop potentially useful options if they're just going to be dropped in the network.

1.2. 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.

2. Requirements for different extension headers

In this section we list the requirements for packets with each of the standard IPv6 extension headers: Hop-by-Hop Options, Encapsulating Security Payload, Routing Header, Destination Options, Fragment Header, and Authentication Header.

2.1. Hop-by-Hop Options

The Hop-by-Hop Options header is used to carry optional information that may be examined and processed by every node along a packet's delivery path. Since Hop-by-Hop Options are processed by routers in the path that makes accepting packets with Hop-by-Hop Options into a limited domain particularly perilous. For instance, an attacker outside a limited domain could employ Hop-by-Hop Options to target the infrastructure of a limited domain with a Denial of Service Attack. As such, this specification strongly recommends that packets with Hop-by-Hop Options are dropped at limited domain boundaries especially at ingress routers.

Requirements:

* At ingress routers of a limited domain, packets with Hop-by-Hop Options SHOULD be discarded. It is highly RECOMMENDED that packets with Hop-by-Hop Options are not allowed to enter a limited domain from untrusted sources.

* At egress routers of a limited domain, the Hop-by-Options header SHOULD be removed from packets before forwarding per the procedures of [inflrm], or packets with Hop-by-Hop Options SHOULD be discarded. Note that removing the Hop-by-Hop Options header instead of dropping allows Hop-by-Hop options to be productively used as packets traverse the limited domain.

2.2. Routing Header

The Routing header is used by an IPv6 source to list one or more intermediate nodes to be "visited" on the way to a packet's destination. The Routing Header has similar properties as Hop-by-Hop Options in that it is processed by routers in the network and therefore Accepting packets with Routing Headers into a limited domain is similarly perilous. It also is well accepted that Routing Headers are only useful in limited domains. Therefore, this specification strongly recommends that packets with a Routing Header are dropped at limited domain boundaries.

Requirements:

* At ingress routers of a limited domain, packets with Routing Header SHOULD be discarded. It is highly RECOMMENDED that packets with a Routing Header are not allowed to enter a limited domain from untrusted sources.

* At egress routers of a limited domain, packets with Routing Header SHOULD be discarded or the Routing Header SHOULD be removed from the packet before forwarding per the procedures of [inflrm]. Note that removing Routing Header instead of dropping allows a Routing Header to be productively used as the packet traverses the limited domain.

2.3. Encapsulating Security Payload (ESP)

Encapsulating Security Payload (ESP) header is designed to provide a mix of security services in IPv4 and IPv6. At routers, parsing beyond the ESP Header isn't possible and the ESP header is in itself strong security, so packets with ESP aren't particularly problematic to accept into limited domains. In practice, some domains have discarded packets with ESP header since they cannot access the transport layer headers, however that is a matter of policy and is not recommended by this specification.

Requirements:

* At ingress routers of a limited domain, packets with an Encapsulating Security Payload SHOULD be accepted.

* At egress routers of a limited domain, packets with an Encapsulating Security Payload SHOULD be accepted.

2.4. Destination Options

The Destination Options header is used to carry optional information that need be examined only by a packet's destination node(s). Destination Options do not directly target routers so accepting packets with them isn't quite as perilous as accepting packets with Hop-by-Hop Options, however they still can have adverse effect in forwarding as described in [RFC9098].

The bigger problem with Destinations Options is that they are sent as plain text (unless they're encapsulated by Encapsulating Security Payload header). Sending plain text over the Internet is an obvious security risk, therefore it is highly recommended that end-to-end information be encrypted when sending over the Internet. Encryption can be accomplished by setting Destination Options after an Encapsulating Security Payload header, or by encapsulating the end to end information in a transport layer encryption header.

Requirements:

* At ingress routers of a limited domain, packets with a Destination Options Header that not encapsulated in an Encapsulating Security Payload header SHOULD be discarded.

* At egress routers of a limited domain, packets with a Destination Options Header that not encapsulated in an Encapsulating Security Payload header SHOULD be discarded. It is likely that packets with Destination Options will be dropped in the Internet anyway so dropping at the limited domain egress router saves unnecessary work.

2.5. Fragment Header

The Fragment header is used by an IPv6 source to send a packet larger then would fit in the path MTU to its destination. The fragility of fragmentation is well documented [RFC8900]. The Fragment header also exposes receiving hosts to Denial of Service attacks. Furthermore, routers cannot access port transport layers in fragments to perform common filtering on port numbers so they often drop packets with the Fragment Header present. Given the issues with the Fragment Header

this specification recommends that packets with a Fragment Header be discarded at limited domain boundaries.

Requirements:

* At ingress routers of a limited domain, packets with a Fragment Headers SHOULD be dropped.

* At egress routers of a limited domain, packets with a Fragment Headers SHOULD be dropped.

2.6. Authentication Header

The IP Authentication Header (AH) is used to provide connectionless integrity and data origin authentication for IP datagrams and to provide protection against replays. The IP Authentication Header has seen little deployment as the functionality is a subset of the ESP header. This specification recommends that packets with an Authentication Header be discarded at limited domain boundaries.

Requirements:

* At ingress routers of a limited domain, packets with an Authentication Header SHOULD be dropped.

* At egress routers of a limited domain, packets with a Authentication Header SHOULD be dropped.

2.7. Sending ICMP errors

If an egress router of a limited domain discards a packet because it disallows an extension header per this specification, the router MAY send an ICMP error message with type of "Parameter Problem" (type 4) and code of "Unrecognized Next Header type encountered by intermediate node" (code 5) per [RFC8883].

If an egress router of a limited domain discards a packet because it disallows an extension Header per this specification, the router SHOULD NOT send an ICMP error. However, if it does send an ICMP error then the router SHOULD send an ICMP error message with type of "Parameter Problem" (type 4) and code of "Unrecognized Next Header type encountered by intermediate node" (code 5) per [RFC8883].

The difference in recommended behavior between ingress and egress routers is that in the egress router case the source host and path to it are in the limited domain so they may be trusted, whereas in the ingress router case the source host and path to it are outside the limited domain and probably not trusted.

3. Requirements for extension headers in limited domains

Within a limited domain it is expected that extension headers can be freely used and managed as necessary. Thus the requirement is that routers in limited domain SHOULD forward packets transparently, independent of the IP protocol type.

4. Acknowledgments

The author would like to thank Brian Carpenter, Justim Iurman and Ole Troan for their comments and suggestions that improved this document.

5. IANA Considerations

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

6. Security Considerations

Security issues with IPv6 extension headers are well known and have been documented in several places including [RFC6398]. By virtue of deprecating the use of extension headers these security issues are no longer relevant. Otherwise, this document does not introduce any new security concerns.

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, <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>.

[RFC8883] Herbert, T., "ICMPv6 Errors for Discarding Packets Due to Processing Limits", RFC 8883, DOI 10.17487/RFC8883, September 2020, <https://www.rfc-editor.org/info/rfc8883>.

7.2. Informative References

[APNIC] Huston, G., "IPv6 Extension headers revisited", October 2022, <https://blog.apnic.net/2022/10/13/ipv6-extension-headers-revisited/>.

[Cus23a] Custura, A. and G. Fairhurst, "Internet Measurements: IPv6 Extension Header Edition", IEPG, IETF-116 , March 2023, <http://www.iepg.org/2023-03-26-ietf116/eh.pdf>.

[Cus23b] Custura, A., Secchi, R., Boswell, E., and G. Fairhurst, "Is it possible to extend IPv6?", Computer Communications X, October 2023, <https://www.sciencedirect.com/science/article/pii/S0140366423003705>.

[eh-limits] Herbert, B., "Limits on Sending and Processing IPv6 Extension Headers", February 2025, <https://www.ietf.org/archive/id/draft-ietf-6man-eh-limits-19.txt>.

[inflrm] Herbert, T., "Infight Removal of IPv6 Hop-by-Hop and Routing Headers", February 2024, <https://www.ietf.org/archive/id/draft-herbert-eh-inflight-removal-04.txt>.

[RFC4942] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/Co-existence Security Considerations", RFC 4942, DOI 10.17487/RFC4942, September 2007, <https://www.rfc-editor.org/info/rfc4942>.

[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>.

[RFC7112] Gont, F., Manral, V., and R. Bonica, "Implications of Oversized IPv6 Header Chains", RFC 7112, DOI 10.17487/RFC7112, January 2014, <https://www.rfc-editor.org/info/rfc7112>.

[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>.

[RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <https://www.rfc-editor.org/info/rfc8799>.

[RFC8900] Bonica, R., Baker, F., Huston, G., Hinden, R., Troan, O., and F. Gont, "IP Fragmentation Considered Fragile", BCP 230, RFC 8900, DOI 10.17487/RFC8900, September 2020, <https://www.rfc-editor.org/info/rfc8900>.

[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>.

Author’s Address

Tom Herbert
XDPnet
Los Gatos, CA,
United States of America
Email: tom@herbertland.com