

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 5 December 2025

F. L. Templin, Ed.
Boeing Research & Technology
3 June 2025

IPv6 Extended Fragment Header for IPv4
draft-templin-intarea-ipid-ext2-10

Abstract

The Internet Protocol, version 4 (IPv4) header includes a 16-bit Identification field in all packets, but this length is too small to ensure reassembly integrity even at moderate data rates in modern networks. Even for Internet Protocol, version 6 (IPv6), the 32-bit Identification field included when a Fragment Header is present may be smaller than desired for some applications. This specification addresses these limitations by adapting the IPv6 Extended Fragment Header for IPv4.

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 5 December 2025.

Copyright Notice

Copyright (c) 2025 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
2. Relation to IPv6	3
3. IPv6 Extended Fragment Header (EFH) for IPv4	3
4. Destination Qualification and Path MTU	4
5. IPv4 Flow Label	4
6. Packet Too Big (PTB) Extensions	4
7. Requirements	5
8. Implementation Status	5
9. IANA Considerations	5
10. Security Considerations	5
11. Acknowledgements	5
12. References	5
12.1. Normative References	6
12.2. Informative References	6
Appendix A. Change Log	7
Author's Address	8

1. Introduction

The Internet Protocol, version 4 (IPv4) header includes a 16-bit Identification in all packets [RFC0791], but this length is too small to ensure reassembly integrity even at moderate data rates in modern networks [RFC4963][RFC6864][RFC8900]. This specification adapts the IPv6 Extended Fragment Header (EFH) [I-D.templin-6man-ipid-ext2] for Identification extension and to support an alternate fragmentation and reassembly service for IPv4.

IPv4 packets that include the IPv6 EFH engage a "deep packet fragmentation" service that supports Identification, fragmentation and reassembly independently of any IPv4 header level services. This may be useful for networks that engage fragmentation and reassembly at extreme data rates, or for cases when advanced packet Identification uniqueness assurance is critical.

2. Relation to IPv6

Protocol extensions intended for IPv6 can often be applied in similar fashion as for IPv4 (and vice-versa). The terminology used and the motivation for extending the Identification field for IPv4 is the same as for the IPv6 Extended Fragment Header (EFH) as specified in [I-D.templin-6man-ipid-ext2]. All normative aspects of the IPv6 specification that can be applied for IPv4 apply also to this document.

3. IPv6 Extended Fragment Header (EFH) for IPv4

IPv4 end systems, intermediate systems and routers by default do not recognize the IP protocol numbers for IPv6 extension headers as these are typically used to support only IPv6 operations. However, implementations of this specification are required to recognize IP protocol number 60 (IPv6 Destination Options header per [RFC8200]) as an applicable extension for IPv4.

Implementations of this specification also recognize the IPv6 EFH Option [I-D.templin-6man-ipid-ext2] when it appears in an IPv6 Destination Options Header following the IPv4 header. Requirements for encapsulation of extension headers in IPv4 packets are introduced and discussed in [I-D.herbert-ipv4-eh]. Recommendations for IPv6 extension header limits are found in [I-D.ietf-6man-eh-limits].

IPv4 sources insert an IPv6 Destination Option Header with an EFH option in an extension header chain beginning immediately after the IPv4 header (plus options) and ending immediately before the upper layer protocol header, e.g., TCP, UDP, etc. The source then increments the IPv4 Total Length by the length of the extension headers, and sets the IPv4 Protocol field to the protocol number of the first extension header. The source then sets the IPv6 Destination Options Next Header field to the protocol number of the next extension header or the upper layer protocol number if there are no further extensions.

The IPv4 source then applies EFH fragmentation if necessary the same as for the IPv6 procedures specified in [I-D.templin-6man-ipid-ext2]. This will produce a sequence of fragment packets each containing a copy of the IPv4 header followed by any Per-Fragment headers up to and including the IPv6 Destination Options Header with EFH Option followed by the fragment payload. The IPv4 source then forwards the fragment packets toward the final destination which processes them only if all nodes in the path pass IPv4 packets that include the IPv6 Destination Options header.

Intermediate systems and IPv4 routers on the path forward the fragment packets if the next hop link MTU is sufficient. A router may perform IPv4 fragmentation if a fragment packet is too large and the Don't Fragment (DF) flag is 0; otherwise, the router drops the packet and returns an ICMPv4 Fragmentation Needed message.

When the fragment packets arrive at the IPv4 destination, it performs IPv4 reassembly if necessary followed by EFH reassembly under the same conditions specified for the IPv6 EFH in [I-D.templin-6man-ipid-ext2].

4. Destination Qualification and Path MTU

IPv4 intermediate systems, routers and destinations that do not recognize the IPv6 Destination Options Header with EFH Option appearing after the IPv4 header unconditionally drop the packet and SHOULD return an "ICMPv4 Destination Unreachable - Protocol Unreachable" message per [RFC0792].

The source can therefore test whether the path up to and including the destination accepts the IPv6 Destination Options Header and EFH Option by occasionally sending "probe" packets of a given size that include them. If the source receives an acknowledgement, it has assurance that the destination recognizes the protocol and that intermediate systems and routers at least forward the protocol messages without dropping; the source can instead consider receipt of an ICMPv4 Destination Unreachable (Protocol Unreachable) as an indication that a node in the path rejects the protocol. The source should occasionally re-probe each destination in case routing redirects a flow to a different anycast destination.

5. IPv4 Flow Label

Destinations that return Fragmentation Report (FR) messages per [I-D.templin-6man-ipid-ext2] set the Flow Label field to the value included by the source of the packet that elicited the FR (an IPv4 Flow Label option appears in [I-D.dreibholz-ipv4-flowlabel]). When the source does not include an explicit Flow Label, the destination sets the field to 0.

6. Packet Too Big (PTB) Extensions

[I-D.templin-6man-ipid-ext2] specifies a new "Packet Too Big (PTB)" ICMP message Type plus Code values for PTB Soft Errors. Intermediate systems and destination end systems return ICMPv4 PTB Soft Errors to the source under the same conditions specified for IPv6.

7. Requirements

All nodes that process IPv4 packets with an IPv6 Destination Options Header including the EFH Option observe the requirements found in [I-D.templin-6man-ipid-ext2] in addition to the requirements found in this section.

Sources SHOULD set DF to 0 and include a suitable IPv4 ID value for packets that include fragment payloads based on the 1024 octet minimum non-final EFH fragment size. Sources MUST set DF to 1 and include any IPv4 ID value for all others. An updated specification of the IPv4 ID field is found in [RFC6864].

Sources MUST include at most one EFH in each IPv4 packet. Intermediate systems and destinations SHOULD silently drop packets with multiples.

Destinations that accept flows using EFH Options MUST configure an EMTU_R of 65535 octets or larger.

8. Implementation Status

In progress.

9. IANA Considerations

This document has no requirements for IANA.

10. Security Considerations

All aspects of IP security apply equally to this document, which does not introduce any new vulnerabilities. Moreover, when employed correctly the mechanisms in this document robustly address known IPv4 reassembly integrity concerns [RFC4963] [RFC8900] and also provide an advanced degree of packet Identification uniqueness assurance.

All other security aspects of the IPv6 Extended Fragment Header per [I-D.templin-6man-ipid-ext2] apply also to its use in IPv4.

11. Acknowledgements

This work was inspired by continued DTN performance studies. Amanda Baber, Tom Herbert, Bob Hinden and Eric Vyncke offered useful insights that helped improve the document.

Honoring life, liberty and the pursuit of happiness.

12. References

12.1. Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981, <<https://www.rfc-editor.org/info/rfc792>>.
- [RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, DOI 10.17487/RFC1191, November 1990, <<https://www.rfc-editor.org/info/rfc1191>>.
- [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>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, RFC 4443, DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [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>>.
- [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>>.

12.2. Informative References

- [I-D.dreibholz-ipv4-flowlabel] Dreibholz, T., "An IPv4 Flowlabel Option", Work in Progress, Internet-Draft, draft-dreibholz-ipv4-flowlabel-41, 30 March 2025, <<https://datatracker.ietf.org/doc/html/draft-dreibholz-ipv4-flowlabel-41>>.

[I-D.herbert-ipv4-eh]

Herbert, T., "IPv4 Extension Headers and Flow Label", Work in Progress, Internet-Draft, draft-herbert-ipv4-eh-03, 22 February 2024, <<https://datatracker.ietf.org/doc/html/draft-herbert-ipv4-eh-03>>.

[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-19, 27 February 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-6man-eh-limits-19>>.

[I-D.templin-6man-ipid-ext2]

Templin, F. and T. Herbert, "IPv6 Extended Fragment Header (EFH)", Work in Progress, Internet-Draft, draft-templin-6man-ipid-ext2-13, 19 May 2025, <<https://datatracker.ietf.org/doc/html/draft-templin-6man-ipid-ext2-13>>.

[RFC4963] Heffner, J., Mathis, M., and B. Chandler, "IPv4 Reassembly Errors at High Data Rates", RFC 4963, DOI 10.17487/RFC4963, July 2007, <<https://www.rfc-editor.org/info/rfc4963>>.

[RFC6864] Touch, J., "Updated Specification of the IPv4 ID Field", RFC 6864, DOI 10.17487/RFC6864, February 2013, <<https://www.rfc-editor.org/info/rfc6864>>.

[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>>.

Appendix A. Change Log

<< RFC Editor - remove prior to publication >>

Differences from version -09 to version -10:

- * Cited draft on IPv6 Flow Labels for IPv4 packets.

Differences from version -07 to version -09:

- * Clarified IPv4 DF bit setting requirements.
- * Included section on IPv4 Flow Label.

- * IPv4 router fragmentation considerations.

Differences from version -06 to version -07:

- * Now using normal (extended) ICMPv4 messages for IPv4 PTB instead of OMNI-encapsulated ICMPv6.

Differences from version -05 to version -06:

- * Removed reference to RFC9268.
- * Clarified setting of DF bit.

Differences from earlier versions:

- * First draft publication.

Author's Address

Fred L. Templin (editor)
Boeing Research & Technology
P.O. Box 3707
Seattle, WA 98124
United States of America
Email: fltemplin@acm.org