6man Working Group
Internet-Draft
Updates: RFC 2460 (if approved)
Intended status: Standards Track
Expires: December 22, 2013

R. Bonica
Juniper Networks
W. Kumari
Google, Inc.

June 20, 2013

IPv6 Fragment Header Deprecated
draft-bonica-6man-frag-deprecate-00

Abstract

This memo deprecates the IPv6 Fragment Header. It provides reasons for deprecation and updates RFC 2460.

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 RFC 2119 [RFC2119].

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 December 22, 2013.

Copyright Notice

Table of Contents

1. Introduction

Each link on the Internet is characterized by a Maximum Transmission Unit (MTU). A link's MTU represents the maximum packet size that can be conveyed over the link, without fragmentation. MTU is a unidirectional metric. A bidirectional link may be characterized by one MTU in the forward direction and another MTU in the reverse direction. IPv6 [RFC2460] requires that every link in the Internet have an MTU of 1280 octets or greater. On any link that cannot convey a 1280-octet packet in one piece, link-specific fragmentation and reassembly must be provided at a layer below IPv6. Therefore, the PMTU between any two IPv6 nodes is 1280 bytes or greater.

Likewise, for any given source node, the path to a particular destination node is characterized by a path MTU (PMTU). At a given source, the PMTU associated with a destination is equal to the minimum MTU of all of the links that contribute to the path between the source and the destination.

[RFC2460] strongly recommends that IPv6 nodes implement Path MTU Discovery (PMTUD) [RFC1981], in order to discover and take advantage of PMTUs greater than 1280 octets. However, a minimal IPv6 implementation (e.g., in a boot ROM) may simply restrict itself to sending packets no larger than 1280 octets, and omit implementation of PMTUD.

In order to send a packet larger than a path's MTU, a node may use the IPv6 Fragment header to fragment the packet at the source and have it reassembled at the destination(s). However, the use of such fragmentation is discouraged in any application that is able to adjust its packets to fit the measured path MTU (i.e., down to 1280 octets).

In IPv6, a packet can be fragmented only by the host that originates it. This constitutes a departure from the IPv4 [RFC0791] fragmentation strategy, in which a packet can be fragmented by its originator or by any router that it traverses en route to its destination.

This memo deprecates the IPv6 Fragment Header. It provides reasons for deprecation and updates [RFC2460].

2. Case For Deprecation

This section presents a case for deprecating the IPv6 Fragment Header.

2.1. Resource Conservation

Packets that are fragmented at their source need to be reassembled at their destination. [Kent87] points out that the reassembly process is resource intensive. It consumes significant compute and memory resources. While the cited reference refers to IPv4 fragmentation and reassembly, many of its criticisms are equally applicable to IPv6.

By comparison, if a source node were to execute PMTUD procedures, and if applications were to avoid sending datagrams that would result in IP packets that exceed the PMTU, the task of reassembly could be avoided, altogether.

2.2. Fragmentation Is Rare

Today, most popular operating systems implement PMTUD or an extension thereof, called Packetization Layer MTU Discovery (PMTUD) [RFC4821]. Most popular TCP [RFC0793] implementations leverage this technology and restrict their segment size so that IP fragmentation is not required. As a result, IPv6 fragments carrying TCP payload are rarely observed on the Internet.

Likewise, many UDP-based [RFC0768] applications follow the recommendations of [RFC5405]. According to [RFC5405], "an application SHOULD NOT send UDP datagrams that result in IP packets that exceed the MTU of the path to the destination. Consequently, an

application SHOULD either use the path MTU information provided by the IP layer or implement path MTU discovery itself to determine whether the path to a destination will support its desired message size without fragmentation. Applications that do not follow this recommendation to do PMTU discovery SHOULD still avoid sending UDP datagrams that would result in IP packets that exceed the path MTU. Because the actual path MTU is unknown, such applications SHOULD fall back to sending messages that are shorter than the default effective MTU for sending." The effective MTU for IPv6 is 1280 bytes.

Because many UDP-based applications follow the above-quoted recommendation, IPv6 fragments carrying UDP traffic are also rarely observed on the Internet.

2.2.1. UDP-based Applications That Rely on Fragmentation

The following is a list of UDP-based applications that do not follow the recommendation of [RFC5405] and rely in IPv6 fragmentation:

- DNSSEC [RFC4035]

The effectiveness of these protocols may currently be degraded by operator behavior. See Section 2.4 for details.

2.3. Attack Vectors

Security researchers have found and continue to find attack vectors that rely on IP fragmentation. For example, [I-D.ietf-6man-oversized-header-chain] and [I-D.ietf-6man-nd-extension-headers] describe variants of the tiny fragment attack [RFC1858]. In this attack, a packet is crafted so that it can evade stateless firewall filters. The stateless firewall filter matches on fields drawn from the IPv6 header and an upper layer header. However, the packet is fragmented so that the upper layer header, or a significant part of that header, does not appear in the first fragment. Because a stateless firewall cannot parse payload beyond the first fragment, the packet evades detection by the firewall.

Security researcher have also studied reassembly algorithms on popular computing platforms, with the following goals:

- to discover fragility in seldom exercised parts of the IP stack
- to engineer flows that maximize resources consumed by the reassembly process

The Dawn and Rose Attacks [Hollis] are the products of such research.

All of the attack vectors mentioned above can be mitigated with firewalls and increasingly sophisticated reassembly algorithms. However, the continued investment required to mitigate newly discovered vulnerabilities detracts from the cost effectiveness of IPv6 as a networking solution.

2.4. Operator Behavior

For reasons described above, and also articulated in [I-D.taylor-v6ops-fragdrop], many network operators filter all IPv6 fragments. Also, the default behavior of many currently deployed firewalls is to discard IPv6 fragments.

In one recent study [DeBoer], two researchers distributed probes to 423 IPv6 enabled sites. The researchers then tested connectivity between an experimental control center and the probes. They found that during any given trial period, sixty percent of the sites that could be reached with unfragmented packets could also be reached with fragmented packets. The remaining forty percent appeared to be filtering IPv6 fragments

3. Recommendation

This memo deprecates IPv6 fragmentation and the IPv6 fragment header. New application and transport layer protocols MUST NOT send datagrams that result in IPv6 packets exceeding the MTU of the path to the destination. However, legacy applications and transport layer protocols will continue to do so.

New IPv6 host implementations MAY support IPv6 fragmentation and reassembly, but are not required to do so.

Network operators MAY filter IPv6 fragments.

4. IANA Considerations

IANA is requested to mark the Fragment Header for IPv6 (44) as deprecated in the Protocol Numbers registry.

5. Security Considerations

Deprecation of the IPv6 Fragment Header will improve network security by eliminating attacks that rely on fragmentation.

6. Acknowledgements

The author wishes to acknowledge Bob Hinden and Ole Troan for their review and constructive comments.

7. References

7.1. Normative References

[RFC0768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.

[RFC1981] McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6", RFC 1981, August 1996.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.

[RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, March 2007.

[RFC5405] Eggert, L. and G. Fairhurst, "Unicast UDP Usage Guidelines for Application Designers", BCP 145, RFC 5405, November 2008.

7.2. Informative References

[DeBoer] De Boer, M. and J. Bosma, "Discovering Path MTU black holes on the Internet using RIPE Atlas", July 2012, <http://www.nlnetlabs.nl/downloads/publications/pmtu-black-holes-msc-thesis.pdf>.

[Hollis] Hollis, K., "The Rose Attack Explained", , <http://digital.net/~gandalf/Rose_Frag_Attack_Explained.htm>.

[I-D.ietf-6man-nd-extension-headers] Gont, F., "Security Implications of IPv6 Fragmentation with IPv6 Neighbor Discovery", draft-ietf-6man-nd-extension-headers-05 (work in progress), June 2013.

[I-D.ietf-6man-oversized-header-chain]
Gont, F. and V. Manral, "Security and Interoperability Implications of Oversized IPv6 Header Chains", draft-ietf-6man-oversized-header-chain-02 (work in progress), November 2012.

[I-D.ietf-6man-predictable-fragment-id]
Gont, F., "Security Implications of Predictable Fragment Identification Values", draft-ietf-6man-predictable-fragment-id-00 (work in progress), March 2013.

[I-D.taylor-v6ops-fragdrop]
Jaeggli, J., Colitti, L., Kumari, W., Vyncke, E., Kaeo, M., and T. Taylor, "Why Operators Filter Fragments and What It Implies", draft-taylor-v6ops-fragdrop-01 (work in progress), June 2013.

[Kent87] Kent, C. and J. Mogul, "Fragmentation Considered Harmful", In Proc. SIGCOMM '87 Workshop on Frontiers in Computer Communications Technology , August 1987.

[RFC1858] Ziemba, G., Reed, D., and P. Traina, "Security Considerations for IP Fragment Filtering", RFC 1858, October 1995.

[RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.

Authors' Addresses

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, Virginia 20170
USA

Email: rbonica@juniper.net

Warren
Google, Inc.
1600 Amphitheatre Parkway
Mountainview, California 94043
USA

Email: warren@kumari.net