

Internet Engineering Task Force (IETF)
Request for Comments: 7112
Updates: 2460
Category: Standards Track
ISSN: 2070-1721

F. Gont
Huawei Technologies
V. Manral
Ionos Networks
R. Bonica
Juniper Networks
January 2014

Implications of Oversized IPv6 Header Chains

Abstract

The IPv6 specification allows IPv6 Header Chains of an arbitrary size. The specification also allows options that can, in turn, extend each of the headers. In those scenarios in which the IPv6 Header Chain or options are unusually long and packets are fragmented, or scenarios in which the fragment size is very small, the First Fragment of a packet may fail to include the entire IPv6 Header Chain. This document discusses the interoperability and security problems of such traffic, and updates RFC 2460 such that the First Fragment of a packet is required to contain the entire IPv6 Header Chain.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7112>.

Copyright Notice

Copyright (c) 2014 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 (http://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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. Introduction2
- 2. Requirements Language3
- 3. Terminology3
- 4. Motivation4
- 5. Updates to RFC 24605
- 6. IANA Considerations5
- 7. Security Considerations6
- 8. Acknowledgements6
- 9. References7
 - 9.1. Normative References7
 - 9.2. Informative References7

1. Introduction

With IPv6, optional internet-layer information is carried in one or more IPv6 Extension Headers [RFC2460]. Extension Headers are placed between the IPv6 header and the Upper-Layer Header in a packet. The term "Header Chain" refers collectively to the IPv6 header, Extension Headers, and Upper-Layer Header occurring in a packet. In those scenarios in which the IPv6 Header Chain is unusually long and packets are fragmented, or scenarios in which the fragment size is very small, the Header Chain may span multiple fragments.

While IPv4 had a fixed maximum length for the set of all IPv4 options present in a single IPv4 packet, IPv6 does not have any equivalent maximum limit at present. This document updates the set of IPv6 specifications to create an overall limit on the size of the combination of IPv6 options and IPv6 Extension Headers that is allowed in a single IPv6 packet. Namely, it updates RFC 2460 such that the First Fragment of a fragmented datagram is required to contain the entire IPv6 Header Chain.

It should be noted that this requirement does not preclude the use of large payloads but, instead, merely requires that all headers, starting from the IPv6 base header and continuing up to the Upper-Layer Header (e.g., TCP or the like) be present in the First Fragment.

2. 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].

3. Terminology

For the purposes of this document, the terms Extension Header, IPv6 Header Chain, First Fragment, and Upper-Layer Header are used as follows:

Extension Header:

Extension Headers are defined in Section 4 of [RFC2460]. As a result of [RFC7045], [IANA-PROTO] provides a list of assigned Internet Protocol Numbers and designates which of those protocol numbers also represent Extension Headers.

First Fragment:

An IPv6 fragment with Fragment Offset equal to 0.

IPv6 Header Chain:

The IPv6 Header Chain contains an initial IPv6 header, zero or more IPv6 Extension Headers, and optionally, a single Upper-Layer Header. If an Upper-Layer Header is present, it terminates the header chain; otherwise, the "No Next Header" value (Next Header = 59) terminates it.

The first member of the IPv6 Header Chain is always an IPv6 header. For a subsequent header to qualify as a member of the header chain, it must be referenced by the "Next Header" field of the previous member of the header chain. However, if a second IPv6 header appears in the header chain, as is the case when IPv6 is tunneled over IPv6, the second IPv6 header is considered to be an Upper-Layer Header and terminates the header chain. Likewise, if an Encapsulating Security Payload (ESP) header appears in the header chain, it is considered to be an Upper-Layer Header, and it terminates the header chain.

Upper-Layer Header:

In the general case, the Upper-Layer Header is the first member of the header chain that is neither an IPv6 header nor an IPv6 Extension Header. However, if either an ESP header, or a second IPv6 header occur in the header chain, they are considered to be Upper-Layer Headers, and they terminate the header chain.

Neither the upper-layer payload, nor any protocol data following the upper-layer payload, is considered to be part of the IPv6 Header Chain. In a simple example, if the Upper-Layer Header is a TCP header, the TCP payload is not part of the IPv6 Header Chain. In a more complex example, if the Upper-Layer Header is an ESP header, neither the payload data, nor any of the fields that follow the payload data in the ESP header are part of the IPv6 Header Chain.

4. Motivation

Many forwarding devices implement stateless firewalls. A stateless firewall enforces a forwarding policy on a packet-by-packet basis. In order to enforce its forwarding policy, the stateless firewall may need to glean information from both the IPv6 and upper-layer headers.

For example, assume that a stateless firewall discards all traffic received from an interface unless it is destined for a particular TCP port on a particular IPv6 address. When this firewall is presented with a fragmented packet that is destined for a different TCP port, and the entire header chain is contained within the First Fragment, the firewall discards the First Fragment and allows subsequent fragments to pass. Because the First Fragment was discarded, the packet cannot be reassembled at the destination. Inasmuch as the packet cannot be reassembled, the forwarding policy is enforced.

However, when the firewall is presented with a fragmented packet and the header chain spans multiple fragments, the First Fragment does not contain enough information for the firewall to enforce its forwarding policy. Lacking sufficient information, the stateless firewall either forwards or discards that fragment. Regardless of the action that it takes, it may fail to enforce its forwarding policy.

5. Updates to RFC 2460

When a host fragments an IPv6 datagram, it **MUST** include the entire IPv6 Header Chain in the First Fragment.

A host that receives a First Fragment that does not satisfy the above-stated requirement **SHOULD** discard the packet and **SHOULD** send an ICMPv6 error message to the source address of the offending packet (subject to the rules for ICMPv6 errors specified in [RFC4443]). However, for backwards compatibility, implementations **MAY** include a configuration option that allows such fragments to be accepted.

Likewise, an intermediate system (e.g., router or firewall) that receives an IPv6 First Fragment that does not satisfy the above-stated requirement **MAY** discard that packet, and it **MAY** send an ICMPv6 error message to the source address of the offending packet (subject to the rules for ICMPv6 error messages specified in [RFC4443]). Intermediate systems having this capability **SHOULD** support configuration (e.g., enable/disable) of whether or not such packets are dropped by the intermediate system.

If a host or intermediate system discards a First Fragment because it does not satisfy the above-stated requirement and sends an ICMPv6 error message due to the discard, then the ICMPv6 error message **MUST** be Type 4 ("Parameter Problem") and **MUST** use Code 3 ("First Fragment has incomplete IPv6 Header Chain"). The Pointer field contained by the ICMPv6 Parameter Problem message **MUST** be set to zero. The format for the ICMPv6 error message is the same regardless of whether a host or intermediate system originates it.

As a result of the above-mentioned requirement, a packet's header chain length cannot exceed the Path MTU associated with its destination. Hosts discover the Path MTU using procedures such as those defined in [RFC1981] and [RFC4821]. Hosts that do not discover the Path MTU **MUST** limit the IPv6 Header Chain length to 1280 bytes. Limiting the IPv6 Header Chain length to 1280 bytes ensures that the header chain length does not exceed the IPv6 minimum MTU [RFC2460].

6. IANA Considerations

IANA has added the following "Type 4 - Parameter Problem" message to the "Internet Control Message Protocol version 6 (ICMPv6) Parameters" registry:

CODE	NAME/DESCRIPTION
3	IPv6 First Fragment has incomplete IPv6 Header Chain

7. Security Considerations

No new security exposures or issues are raised by this document. This document describes how undesirably fragmented packets can be leveraged to evade stateless packet filtering. Having made that observation, this document updates [RFC2460] so that undesirably fragmented packets are forbidden. Therefore, a security vulnerability is removed.

This specification allows nodes that drop the aforementioned packets to signal such packet drops with ICMPv6 "Parameter Problem, IPv6 First Fragment has incomplete IPv6 header chain" (Type 4, Code 3) error messages.

As with all ICMPv6 error/diagnostic messages, deploying Source Address Forgery Prevention filters helps reduce the chances of an attacker successfully performing a reflection attack by sending forged illegal packets with the victim's/target's IPv6 address as the IPv6 source address of the illegal packet [RFC2827] [RFC3704].

A firewall that performs stateless deep packet inspection (i.e., examines application payload content) might still be unable to correctly process fragmented packets, even if the IPv6 Header Chain is not fragmented.

8. Acknowledgements

The authors would like to thank Ran Atkinson for contributing text and ideas that were incorporated into this document.

The authors would like to thank (in alphabetical order) Ran Atkinson, Fred Baker, Stewart Bryant, Brian Carpenter, Benoit Claise, Dominik Elsbroek, Wes George, Mike Heard, Bill Jouris, Suresh Krishnan, Dave Thaler, Ole Troan, Eric Vyncke, and Peter Yee, for providing valuable comments on earlier versions of this document.

9. References

9.1. Normative References

- [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.
- [RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", RFC 7045, December 2013.

9.2. Informative References

- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [IANA-PROTO] Internet Assigned Numbers Authority, "Protocol Numbers", <<http://www.iana.org/assignments/protocol-numbers>>.

Authors' Addresses

Fernando Gont
Huawei Technologies
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
EMail: fgont@si6networks.com

Vishwas Manral
Ionos Networks
Sunnyvale, CA 94089
US

Phone: 408-447-1497
EMail: vishwas@ionosnetworks.com

Ronald P. Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, VA 20171
US

Phone: 571 250 5819
EMail: rbonica@juniper.net