IPv6 maintenance Working Group (6man)
Internet-Draft
Updates: 2460 (if approved)
Intended status: Standards Track
Expires: March 14, 2014

F. Gont
SI6 Networks / UTN-FRH
V. Manral
Hewlett-Packard Corp.
R. Bonica
Juniper Networks

September 10, 2013

Implications of Oversized IPv6 Header Chains
draft-ietf-6man-oversized-header-chain-07

Abstract

The IPv6 specification allows IPv6 header chains of an arbitrary size. The specification also allows options which 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

Copyright Notice

Table of Contents

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 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 RFC 2119 [RFC2119].

3. Terminology

For the purposes of this document, the terms Extension Header, 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 [I-D.ietf-6man-ext-transmit], [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 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 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 header chain. In a simple example, if the upper-layer header is a TCP header, the TCP payload is not part of the 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 header chain.

4. Motivation

Many forwarding devices implement stateless firewalls. A stateless firewall enforces a forwarding policy on 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 destined for a particular TCP port on a particular IPv6 address. When this firewall is presented with a fragmented packet, 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. Insomuch 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 a IPv6 datagram, it MUST include the entire header chain in the first fragment.

A host that receives a first-fragment that does not satisfy the above-stated requirement SHOULD discard that packet, and also MAY send an ICMPv6 error message to the source address of the offending packet (subject to the rules for ICMPv6 errors specified in [RFC4443]).

Likewise, an intermediate system (e.g. router, firewall) that receives an IPv6 first-fragment that does not satisfy the above-stated requirements MAY discard that packet, and 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 such packets are dropped or not by the intermediate system.

If a host or intermediate system discards a first-fragment because it does not satisfy the above-stated requirements, 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 TBD ("First-fragment has incomplete IPv6 Header Chain"). The Pointer field contained by the ICMPv6 Parameter Problem message MUST be set to zero.

As a result of the above mentioned requirements, a packet's header chain length cannot exceed the Path MTU associated with its destination. Hosts MAY discover the Path MTU, using procedures such as those defined in [RFC1981] and [RFC4821]. However, if a host does not discover the Path MTU, it MUST limit the header chain length to 1280 bytes. Limiting the header chain length to 1280 bytes ensures that the header chain length does not exceed the IPv6 minimum MTU.

6. IANA Considerations

IANA is requested to add a the following entry to the "Reason Code" registry for ICMPv6 "Type 4 - Parameter Problem" messages:

CODE	NAME/DESCRIPTION
TBD	IPv6 first-fragment has incomplete IPv6 header chain

7. Security Considerations

This document describes how improperly-fragmented packets can prevent traditional stateless packet filtering.

This document updates RFC 2460 such that those packets are forbidden, thus enabling stateless packet filtering for IPv6.

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 TBD) 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/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 of this document 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, Brian Carpenter, Dominik Elsbroek, Wes George, Mike Heard, Bill Jouris, Suresh Krishnan, Dave Thaler, Ole Troan, and Eric Vyncke, 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.

[I-D.ietf-6man-ext-transmit] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", draft-ietf-6man-ext-transmit-03 (work in progress), August 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", February 2013, <http://www.iana.org/assignments/protocol-numbers/protocol-numbers.txt>.

Authors' Addresses

Fernando Gont
SI6 Networks / UTN-FRH
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
Email: fgont@si6networks.com
URI: http://www.si6networks.com

Vishwas Manral
Hewlett-Packard Corp.
191111 Pruneridge Ave.
Cupertino, CA 95014
US

Phone: 408-447-1497
Email: vishwas.manral@hp.com
URI:

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

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