IPv6 maintenance Working Group (6man)
Internet-Draft
Updates: 2460 (if approved)
Intended status: Standards Track
Expires: May 9, 2013
November 5, 2012

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

Security and Interoperability Implications of Oversized IPv6 Header Chains
draft-ietf-6man-oversized-header-chain-02

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, IPv6 options are carried inside one or more IPv6 Extension Headers [RFC2460]. A sequence of more than one IPv6 Extension Headers in a row is commonly called an "IPv6 Header Chain". 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 first fragment of a packet may fail to include the entire IPv6 header chain.

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 e.g. IPv6 jumbo 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. Terminology

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].

IPv6 Extension Headers:
Any Extension Headers as described in Section 4 of [RFC2460], and specified in [RFC2460] or any subsequent documents.

Entire IPv6 header chain:
All protocol headers starting from the fixed IPv6 header up to (and including) the upper layer protocol header (TCP, UDP, etc. -- assuming there is one of those), including any intermediate IPv6 extension headers.

Note: If there is an upper layer header, only the header (and not its payload) are considered part of the "entire IPv6 header chain". For example, if the upper layer protocol is TCP, only the TCP header (and not its possible data bytes) should be considered part of the "entire IPv6 header chain".

3. Interoperability Implications of Oversized IPv6 Header Chains

Some transition technologies, such as NAT64 [RFC6146], might need to have access to the entire IPv6 header chain in order to associate an incoming IPv6 packet with an ongoing "session".

> For instance, Section 3.4 of [RFC6146] states that "The NAT64 MAY require that the UDP, TCP, or ICMP header be completely contained within the fragment that contains fragment offset equal to zero".

Failure to include the entire IPv6 header chain in the first-fragment might cause the translation function to fail, with the corresponding packets being dropped.

4. Forwarding Implications of Oversized IPv6 Header Chains

A lot of the switches and Routers in the internet do hardware based forwarding. To be able to achieve a level of throughput, there is a fixed maximum number of clock cycles dedicated to each packet. However with the use of unlimited options and header interleaving, larger packets with a lot of interleaving might have to be forwarded to the software. This is one reason why the maximum size of valid packets with interleaved headers needs to be limited.

5. Security Implications of Oversized IPv6 Header Chains

Most firewalls enforce their filtering policy based on the following parameters:

- Source IP address
- Destination IP address
- Protocol type (e.g. ICMPv6, TCP, UDP, SCTP)
- Transport-layer Source Port number
- Transport-layer Destination Port number

Some firewalls reassemble fragmented packets before applying a filtering policy, and thus always have the aforementioned information available when deciding whether to allow or block a packet. However, other stateless firewalls (generally prevalent on small/ home office equipment) do not reassemble fragmented traffic, and hence have to enforce their filtering policy based on the information contained in the received fragment, as opposed to the information contained in the reassembled datagram.

When presented with fragmented traffic, many of such firewalls typically enforce their policy only on the first fragment of a packet, based on the assumption that if the first fragment is dropped, reassembly of the corresponding datagram will fail, and thus such datagram will be effectively blocked. However, if the first fragment fails to include the entire IPv6 header chain, they might have no alternative other than "blindly" allowing or blocking the corresponding fragment. If they blindly allow the packet, then the firewall can be easily circumvented by intentionally sending fragmented packets that fail to include the entire IPv6 header chain in the first fragment. On the other hand, first-fragments that fail to include the entire IPv6 header chain have never been formally deprecated and thus, in theory, might be legitimately generated.

6. Updating RFC 2460

If an IPv6 packet is fragmented, the first fragment of that IPv6 packet (i.e., the fragment having a Fragment Offset of 0) MUST contain the entire IPv6 header chain.

A host that receives an IPv6 first-fragment that does not contain the entire IPv6 header chain SHOULD drop that packet, and also MAY send an ICMPv6 error message to the (claimed) source address (subject to the sending rules for ICMPv6 errors specified in [RFC4443]).

An intermediate system (e.g. router, firewall) that receives an IPv6 first-fragment that does not contain the entire IPv6 header chain MAY drop that packet, and MAY send an ICMPv6 error message to the (claimed) source address (subject to the sending 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 drops an IPv6 first-fragment because it does not contain the entire IPv6 Header Chain, and sends an ICMPv6 error message due to that packet drop, then the ICMPv6 error message MUST be Type 4 ("Parameter Problem") and MUST use Code 3 ("First-fragment has incomplete IPv6 Header Chain").

Implementations SHOULD support configuration of whether an ICMPv6 error/diagnostic message is sent when such packet drops occur. Implementations might consider providing not only an enable/disable configuration, but also other settings (e.g. rate-limit the sending of this kind of ICMPv6 error message).

Sending this ICMPv6 error message when such packets are dropped can be very helpful in diagnosing operational IPv6 network problems, for example if recursive tunnels or certain link technologies have reduced the end-to-end MTU from larger more common values. However, such ICMPv6 messages also might be operationally problematic, for example if an adversary forges the source address on IPv6 first-fragment packets that do NOT contain the entire IPv6 Header Chain. So configurability about sending these ICMPv6 error messages is very important to network operators for this situation.

7. IANA Considerations

IANA is requested that the "Reason Code" registry for ICMPv6 "Type 4 - Parameter Problem" messages be updated as follows:

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

8. Security Considerations

This document describes the interoperability and security implications of IPv6 packets or first-fragments that fail to include the entire IPv6 header chain. The security implications include the possibility of an attacker evading network security controls such as firewalls and Network Intrusion Detection Systems (NIDS) [CPNI-IPv6].

This document updates RFC 2460 such that those packets are forbidden, thus preventing the aforementioned issues.

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/target's IPv6 address as the IPv6 Source Address of the illegal packet [RFC2827] [RFC3704].

9. 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, Dominik Elsbroek, Bill Jouris, Suresh Krishnan, Dave Thaler, and Eric Vyncke, for providing valuable comments on earlier versions of this document.

10. References

10.1. Normative References

[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.

10.2. 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.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

[CPNI-IPv6] Gont, F., "Security Assessment of the Internet Protocol version 6 (IPv6)", UK Centre for the Protection of National Infrastructure, (available on request).

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: