IPv6 maintenance Working Group (6man)
Internet-Draft
Updates: 2460 (if approved)
Intended status: Standards Track
Expires: February 20, 2015

F. Gont
SI6 Networks / UTN-FRH
W. Liu
Huawei Technologies

August 19, 2014

Deprecating the Generation of IPv6 Atomic Fragments
draft-gont-6man-deprecate-atomfrag-generation-00

Abstract

The core IPv6 specification requires that when a host receives an ICMPv6 "Packet Too Big" message reporting a "Next-Hop MTU" smaller than 1280, the host includes a Fragment Header in all subsequent packets sent to that destination, without reducing the assumed Path-MTU. The simplicity with which ICMPv6 "Packet Too Big" messages can be forged, coupled with the widespread filtering of IPv6 fragments, results in an attack vector that can be leveraged for Denial of Service purposes. This document briefly discusses the aforementioned attack vector, and formally deprecates the generation of IPv6 atomic fragments, such that the aforementioned attack vector is eliminated.

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 February 20, 2015.

Copyright Notice

Table of Contents

1. Introduction

[RFC2460] specifies the IPv6 fragmentation mechanism, which allows IPv6 packets to be fragmented into smaller pieces such that they fit in the Path-MTU to the intended destination(s).

Section 5 of [RFC2460] states that, when a host receives an ICMPv6 "Packet Too Big" message [RFC4443] advertising a "Next-Hop MTU" smaller than 1280 (the minimum IPv6 MTU), the host is not required to reduce the assumed Path-MTU, but must simply include a Fragment Header in all subsequent packets sent to that destination. The resulting packets will thus *not* be actually fragmented into several pieces, but rather just include a Fragment Header with both the "Fragment Offset" and the "M" flag set to 0 (we refer to these packets as "atomic fragments"). As required by [RFC6946], these atomic fragments are essentially processed by the destination host as non-fragment traffic (since there are not really any fragments to be reassembled). IPv6/IPv4 translators will typically employ the Fragment Identification information found in the Fragment Header to select an appropriate Fragment Identification value for the resulting IPv4 fragments.

While atomic fragments might seem rather benign, there are scenarios in which the generation of IPv6 atomic fragments can introduce an attack vector that can be exploited for denial of service purposes. Since there are concrete security implications arising from the

generation of IPv6 atomic fragments, and there is no real gain in generating IPv6 atomic fragments (as opposed to e.g. having IPv6/IPv4 translators generate a Fragment Identification value themselves), this document formally updates [RFC2460], forbidding the generation of IPv6 atomic fragments, such that the aforementioned attack vector is eliminated.

Section 3 describes some possible attack scenarios. Section 5 provides additional considerations regarding the usefulness of generating IPv6 atomic fragments. Section 4 formally updates RFC2460 such that this attack vector is eliminated.

2. Terminology

IPv6 atomic fragments
: IPv6 packets that contain a Fragment Header with the Fragment Offset set to 0 and the M flag set to 0 (as defined by [RFC6946]).

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. Denial of Service (DoS) attack vector

Let us assume that Host A is communicating with Server B, and that, as a result of the widespread filtering of IPv6 packets with extension headers (including fragmentation) [I-D.gont-v6ops-ipv6-ehs-in-real-world], some intermediate node filters fragments between Host A and Server B. If an attacker sends a forged ICMPv6 "Packet Too Big" (PTB) error message to server B, reporting a Next-Hop MTU smaller than 1280, this will trigger the generation of IPv6 atomic fragments from that moment on (as required by [RFC2460]). When server server B starts sending IPv6 atomic fragments (in response to the received ICMPv6 PTB), these packets will be dropped, since we previously noted that packets with IPv6 EHs were being dropped between Host A and Server B. Thus, this situation

will result in a Denial of Service (DoS) scenario.

Another possible scenario is that in which two BGP peers are employing IPv6 transport, and they implement ACLs to drop IPv6 fragments (to avoid control-plane attacks). If the aforementioned BGP peers drop IPv6 fragments but still honor received ICMPv6 Packet Too Big error messages, an attacker could easily attack the peering session by simply sending an ICMPv6 PTB message with a reported MTU smaller than 1280 bytes. Once the attack packet has been fired, it will be the aforementioned routers themselves the ones dropping their own traffic.

The aforementioned attack vector is exacerbated by the following factors:

- The attacker does not need to forge the IPv6 Source Address of his attack packets. Hence, deployment of simple BCP38 filters will not help as a counter-measure.

- Only the IPv6 addresses of the IPv6 packet embedded in the ICMPv6 payload need to be forged. While one could envision filtering devices enforcing BCP38-style filters on the ICMPv6 payload, the use of extension (by the attacker) could make this difficult, if at all possible.

- Many implementations fail to perform validation checks on the received ICMPv6 error messages, as recommended in Section 5.2 of [RFC4443] and documented in [RFC5927]. It should be noted that in some cases, such as when an ICMPv6 error message has (supposedly) been elicited by a connection-less transport protocol (or some other connection-less protocol being encapsulated in IPv6), it may be virtually impossible to perform validation checks on the received ICMPv6 error messages. And, because of IPv6 extension headers, the ICMPv6 payload might not even contain any useful information on which to perform validation checks.

- Upon receipt of one of the aforementioned ICMPv6 "Packet Too Big" error messages, the Destination Cache [RFC4861] is usually updated to reflect that any subsequent packets to such destination should include a Fragment Header. This means that a single ICMPv6 "Packet Too Big" error message might affect multiple communication

instances (e.g., TCP connections) with such destination.

4. Updating RFC2460

The following text from Section 5 of [RFC2460]:

> "In response to an IPv6 packet that is sent to an IPv4 destination (i.e., a packet that undergoes translation from IPv6 to IPv4), the originating IPv6 node may receive an ICMP Packet Too Big message reporting a Next-Hop MTU less than 1280. In that case, the IPv6 node is not required to reduce the size of subsequent packets to less than 1280, but must include a Fragment header in those packets so that the IPv6-to-IPv4 translating router can obtain a suitable Identification value to use in resulting IPv4 fragments. Note that this means the payload may have to be reduced to 1232 octets (1280 minus 40 for the IPv6 header and 8 for the Fragment header), and smaller still if additional extension headers are used."

is formally replaced with:

> "IPv6 nodes MUST discard ICMPv6 Packet Too Big error messages that report a Next-Hop MTU smaller than 1280 bytes (the minimum IPv6 MTU)."

5. Additional Considerations

Besides the security assessment provided in Section 3, it is interesting to evaluate if there is any gain in generating IPv6 atomic fragments (to provide for Fragment Identification value) as opposed to just let IPv6/IPv4 translators select an appropriate IPv4 Fragment Identification value.

After some analysis, one can conclude that, if anything, an IPv6/IPv4 translator is in a much better position to select an appropriate Fragment Identification value for the packet that are to be translated from the IPv6 to the IPv4 world. For instance, an IPv6 node will generate Fragment Identification values without any knowledge of the Fragment ID values being generated by other IPv6 nodes employing the translator. Thus, an IPv6/IPv4 translator is in a much better position to generate Fragment IDs that will not result

in collisions (i.e., that will not be reused for the same tuple {Source Address, Destination Address}.

6. IANA Considerations

There are no IANA registries within this document. The RFC-Editor can remove this section before publication of this document as an RFC.

7. Security Considerations

This document describes a Denial of Service (DoS) attack vector that leverages the widespread filtering of IPv6 fragments in the public Internet by means of ICMPv6 PTB error messages. Additionally, it formally updates [RFC2460] such that this attack vector is eliminated.

8. Acknowledgements

Fernando Gont would like to thank Jan Zorz and Go6 Lab <http://go6lab.si/> for providing access to systems and networks that were employed to produce some of the measurement results presented in this document. Additionally, he would like to thank SixXS <https://www.sixxs.net> for providing IPv6 connectivity.

9. References

9.1. Normative References

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

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

[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.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman,

"Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

9.2. Informative References

[RFC5927] Gont, F., "ICMP Attacks against TCP", RFC 5927, July 2010.

[RFC6946] Gont, F., "Processing of IPv6 "Atomic" Fragments", RFC 6946, May 2013.

[I-D.gont-v6ops-ipv6-ehs-in-real-world] Gont, F., Linkova, J., Chown, T., and W. Will, "IPv6 Extension Headers in the Real World", draft-gont-v6ops-ipv6-ehs-in-real-world-00 (work in progress), August 2014.

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

Will(Shucheng) Liu
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China

Email: liushucheng@huawei.com