

IPv6 maintenance Working Group (6man)
Internet-Draft
Updates: [2460](#) (if approved)
Intended status: Best Current Practice
Expires: November 1, 2014

F. Gont
SI6 Networks / UTN-FRH
April 30, 2014

Security Implications of Predictable Fragment Identification Values
draft-ietf-6man-predictable-fragment-id-01

Abstract

IPv6 specifies the Fragment Header, which is employed for the fragmentation and reassembly mechanisms. The Fragment Header contains an "Identification" field which, together with the IPv6 Source Address and the IPv6 Destination Address of a packet, identifies fragments that correspond to the same original datagram, such that they can be reassembled together at the receiving host. The only requirement for setting the "Identification" value is that it must be different than that employed for any other fragmented packet sent recently with the same Source Address and Destination Address. Some implementations use simple a global counter for setting the Identification field, thus leading to predictable values. This document analyzes the security implications of predictable Identification values, and updates [RFC 2460](#) specifying additional requirements for setting the Identification field of the Fragment Header, such that the aforementioned security implications are mitigated.

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 November 1, 2014.

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.	Introduction	2
2.	Terminology	3
3.	Security Implications of Predictable Fragment Identification values	3
4.	Updating RFC 2460	6
5.	Constraints for the selection of Fragment Identification Values	6
6.	Algorithms for Selecting Fragment Identification Values	7
6.1.	Per-destination counter (initialized to a random value)	7
6.2.	Randomized Identification values	8
6.3.	Hash-based Fragment Identification selection algorithm	9
7.	IANA Considerations	11
8.	Security Considerations	11
9.	Acknowledgements	11
10.	References	11
10.1.	Normative References	11
10.2.	Informative References	12
Appendix A.	Information leakage produced by vulnerable implementations	13
Appendix B.	Survey of Fragment Identification selection algorithms employed by popular IPv6 implementations	15
	Author's Address	16

[1.](#) Introduction

IPv6 specifies the Fragment Header, which is employed for the fragmentation and reassembly mechanisms. The Fragment Header contains an "Identification" field which, together with the IPv6 Source Address and the IPv6 Destination Address of a packet, identifies fragments that correspond to the same original datagram, such that they can be reassembled together at the receiving host.

The only requirement for setting the "Identification" value is that it must be different than that employed for any other fragmented packet sent recently with the same Source Address and Destination Address.

The most trivial algorithm to avoid reusing Fragment Identification values too quickly is to maintain a global counter that is incremented for each fragmented packet that is transmitted. However, this trivial algorithm leads to predictable Identification values, which can be leveraged to performing a variety of attacks.

[Section 3](#) of this document analyzes the security implications of predictable Identification values. [Section 4](#) updates [RFC 2460](#) by adding the requirement that IPv6 Fragment Identification values must not be predictable by an off-path attacker. [Section 5](#) discusses constraints in the possible algorithms for selecting Fragment Identification values. [Section 6](#) specifies a number of algorithms that could be used for generating Identification values. Finally, [Appendix B](#) contains a survey of the Fragment Identification algorithms employed by popular IPv6 implementations.

[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](#)].

[3.](#) Security Implications of Predictable Fragment Identification values

Predictable Identification values result in an information leakage that can be exploited in a number of ways. Among others, they may potentially be exploited to:

- o determine the packet rate at which a given system is transmitting information,
- o perform stealth port scans to a third-party,
- o uncover the rules of a number of firewalls,
- o count the number of systems behind a middle-box,
- o perform Denial of Service (DoS) attacks, or,
- o perform data injection attacks against transport or application protocols

[CPNI-IPv6] contains a detailed analysis of possible vulnerabilities introduced by predictable Fragment Identification values. In summary, their security implications are very similar to those of predictable Identification values in IPv4.

[Sanfilippo1998a] originally pointed out how the IPv4 Identification field could be examined to determine the packet rate at which a given system is transmitting information. Later, [Sanfilippo1998b] describes how a system with such an implementation could be used to perform a stealth port scan to a third (victim) host. [Sanfilippo1999] explains how to exploit this implementation strategy to uncover the rules of a number of firewalls. [Bellovin2002] explains how the IPv4 Identification field can be exploited to count the number of systems behind a NAT. [Fyodor2004] is an entire paper on most (if not all) the ways to exploit the information provided by the Identification field of the IPv4 header (and these results apply in a similar way to IPv6). [Zalewski2003] originally envisioned the exploitation of IP fragmentation for performing data injection attacks against upper-layer protocols. [Herzberg2013] explores the use of IPv4/IPv6 fragmentation and predictable Identification values for performing DNS cache poisoning attacks in great detail. [RFC6274] covers the security implications of the IPv4 case in detail.

One key difference between the IPv4 case and the IPv6 case is that in IPv4 the Identification field is part of the fixed IPv4 header (and thus usually set for all packets), while in IPv6 the Identification field is present only in those packets that carry a Fragment Header. As a result, successful exploitation of the IPv6 Fragment Identification field depends on two different factors:

- o vulnerable IPv6 Fragment Identification generators, and,
- o the ability of an attacker to trigger the use of IPv6 fragmentation for packets sent from/to the victim node

As noted in the previous section, some implementations have been known to use predictable Fragment Identification values. For instance, [Appendix B](#) of this document shows that recent versions of a number of popular IPv6 implementations have employed predictable values for the IPv6 Fragment Identification.

Additionally, we note that [RFC 1981](#) [RFC1981] states that when an ICMPv6 Packet Too Big error message advertising an MTU smaller than 1280 bytes is received, the receiving host is not required to reduce the Path-MTU for the corresponding destination address, but must simply include a Fragment Header in all subsequent packets sent to that destination. This triggers the use of the so-called IPv6

"atomic fragments" [[RFC6946](#)]: IPv6 fragments with a Fragment Offset equal to 0, and the "M" ("More fragments") bit clear.

Thus, an attacker can usually cause a victim host to "fragment" its outgoing packets by sending it a forged ICMPv6 'Packet Too Big' (PTB) error message that advertises a Next-Hop MTU smaller than 1280 bytes.

There are a number of aspects that should be considered, though:

- o All the implementations the author is aware of record the Path-MTU information on a per-destination basis. Thus, an attacker can only cause the victim to enable fragmentation for those packets sent to the Source Address of IPv6 packet embedded in the payload of the ICMPv6 PTB message. However, we note that [Section 5.2 of \[RFC1981\]](#) notes that an implementation could maintain a single system-wide PMTU value to be used for all packets originating from that nodes. Clearly, such an implementations would exacerbate the problem of any attacks based on PMTUD [[RFC5927](#)] or IPv6 fragmentation.
- o If the victim node implements some of the counter-measures for ICMP attacks described in [RFC 5927](#) [[RFC5927](#)], it might be difficult for an attacker to cause the victim node to use fragmentation for its outgoing packets. However, many current implementations fail to enforce these validation checks. For example, Linux 2.6.38-8 does not even require received ICMPv6 error messages to correspond to ongoing communication instances.

Implementations that employ predictable Identification values and also fail to enforce validation checks on ICMPv6 error messages become vulnerable to the same type of attacks that can be exploited with IPv4 fragmentation, discussed earlier in this section.

One possible way in which predictable Identification values could be leveraged for performing a Denial of Service (DoS) attack is as follows: Let us assume that Host A is communicating with Host B, and that an attacker wants to DoS such communication. The attacker would learn the the Identification value currently in use by Host A, possibly by sending any packet that would elicit a fragmented response (e.g., an ICCPMv6 echo request with a large payload). The attacker would then send a forged ICMPv6 Packet Too Big error message to Host A (with the IPv6 Destination Address of the embedded IPv6 packet set to the IPv6 address of a Host B), such that any subsequent packets sent by Host A to Host B include a Fragment Header. Finally, the attacker send forged IPv6 fragments to the Host B, with their IPv6 Source Address set to that of Host A, and Identification values that would result in collisions with the Identification values employed for the legitimate traffic sent by Host A to Host B. If Host

B discards fragments that result in collisions of Identification values (e.g., such fragments overlap, and the host implements [[RFC5722](#)]), the attacker could simply trash the Identification space by sending multiple forged fragments with different Identification values, such that any subsequent packets from Host A to Host B are discarded at Host B as a result of the malicious fragments sent by the attacker.

NOTES:

For example, Linux 2.6.38-10 is vulnerable to the aforementioned issue.

[RFC6946] describes an improved processing of these packets that would eliminate this specific attack vector, at least in the case of TCP connections that employ the Path-MTU Discovery mechanism.

The previous attack scenario is simply included to illustrate the problem of employing predictable fragment Identification values. We note that regardless of the attacker's ability to cause a victim host to employ fragmentation when communicating with third-parties, use of predictable Identification values makes communication flows that employ fragmentation vulnerable to any fragmentation-based attacks.

4. Updating [RFC 2460](#)

Hereby we update [RFC 2460](#) [[RFC2460](#)] as follows:

The Identification value of the Fragment Header MUST NOT be predictable by an off-path attacker.

5. Constraints for the selection of Fragment Identification Values

The "Identification" field of the Fragmentation Header is 32-bits long. However, when translators [[RFC6145](#)] are employed, the "effective" length of the IPv6 Fragment Identification field is 16 bits.

NOTE: [[RFC6145](#)] notes that, when translating in the IPv6-to-IPv4 direction, "if there is a Fragment Header in the IPv6 packet, the last 16 bits of its value MUST be used for the IPv4 identification value". This means that the high-order 16 bits are effectively ignored.

As a result, at least during the IPv6/IPv4 transition/co-existence phase, it is probably safer to assume that only the low-order 16 bits of the IPv6 Fragment Identification are of use to the destination system.

Regarding the selection of Fragment Identification values, the only requirement specified in [[RFC2460](#)] is that the Fragment Identification must be different than that of any other fragmented packet sent recently with the same Source Address and Destination Address. Failure to comply with this requirement could lead to the interoperability problems discussed in [[RFC4963](#)].

From a security standpoint, unpredictable Identification values are desirable. However, this is somewhat at odds with the "re-use" requirements specified in [[RFC2460](#)].

Finally, since Fragment Identification values need to be selected for each outgoing datagram that requires fragmentation, the performance aspect should be considered when choosing an algorithm for the selection of Fragment Identification values.

6. Algorithms for Selecting Fragment Identification Values

This section specifies a number of algorithms that MAY be used for selecting Fragment Identification values.

6.1. Per-destination counter (initialized to a random value)

1. Whenever a packet must be sent with a Fragment Header, the sending host should perform a look-up in the Destinations Cache an entry corresponding to the Destination Address of the packet.
2. If such an entry exists, it contains the last Fragment Identification value used for that Destination. Therefore, such value should be incremented by 1, and used for setting the Fragment Identification value of the outgoing packet. Additionally, the updated value should be recorded in the corresponding entry of the Destination Cache.
3. If such an entry does not exist, it should be created, and the "Identification" value for that destination should be initialized with a random value (e.g., with a pseudorandom number generator), and used for setting the Identification field of the Fragment Header of the outgoing packet.

The advantages of this algorithm are:

- o It is simple to implement, with the only complexity residing in the Pseudo-Random Number Generator (PRNG) used to initialize the "Identification" value contained in each entry of the Destinations Cache.

- o The "Identification" re-use frequency will typically be lower than that achieved by a global counter (when sending traffic to multiple destinations), since this algorithm uses per-destination counters (rather than a single system-wide counter).
- o It has good performance properties (once the corresponding entry in the Destinations Cache has been created, each subsequent "Identification" value simply involves the increment of a counter).

The possible drawbacks of this algorithm are:

- o If as a result of resource management an entry of the Destinations Cache must be removed, the last Fragment Identification value used for that Destination will be lost. Thus, subsequent traffic to that destination would cause that entry to be re-created and re-initialized to random value, thus possibly leading to Fragment Identification "collisions".
- o Since the Fragment Identification values are predictable by the destination host, a vulnerable host might possibly leak to third-parties the Fragment Identification values used by other hosts to send traffic to it (i.e., Host B could leak to Host C the Fragment Identification values that Host A is using to send packets to Host B). [Appendix A](#) describes one possible scenario for such leakage in detail.

6.2. Randomized Identification values

Clearly, use of a Pseudo-Random Number Generator for selecting the Fragment Identification would be desirable from a security standpoint. With such a scheme, the Fragment Identification of each fragmented datagram would be selected as:

```
Identification = random()
```

where "random()" is the PRNG.

The specific properties of such scheme would clearly depend on the specific PRNG algorithm used. For example, some PRNGs may result in higher Fragment Identification reuse frequencies than others, in the same way as some PRNGs may be more expensive (in terms of processing requirements and/or implementation complexity) than others.

Discussion of the properties of possible PRNGs is considered out of the scope of this document. However, we do note that some PRNGs employed in the past by some implementations have been found to be

predictable [[Klein2007](#)]. Please see [[RFC4086](#)] for randomness requirements for security.

6.3. Hash-based Fragment Identification selection algorithm

Another alternative is to implement a hash-based algorithm similar to that specified in [[RFC6056](#)] for the selection of transport port numbers. With such a scheme, the Fragment Identification value of each fragment datagram would be selected with the expression:

$$\text{Identification} = F(\text{Src IP}, \text{Dst IP}, \text{secret1}) + \text{counter}[G(\text{src IP}, \text{Dst Pref}, \text{secret2})]$$

where:

Identification:

Identification value to be used for the fragmented datagram

F():

Hash function

Src IP:

IPv6 Source Address of the datagram to be fragmented

Dst IP:

IPv6 Destination Address of the datagram to be fragmented

secret1:

Secret data unknown to the attacker

counter[]:

System-wide array of 32-bit counters (e.g. with 8K elements or more)

G():

Hash function. May or may not be the same hash function as that used for F()

Dst Pref:

IPv6 "Destination Prefix" of datagram to be fragmented (can be assumed to be the first eight bytes of the Destination Address of such packet). Note: the "Destination Prefix" (rather than Destination Address) is used, such that the ability of an attacker of searching the "increments" space by using multiple addresses of the same subnet is reduced.

secret1:

Secret data unknown to the attacker

NOTE: counter[G(src IP, Dst Pref, secret2)] should be incremented by one each time an Identification value is selected.

The advantages of this algorithm are:

- o The "Identification" re-use frequency will typically be lower than that achieved by a global counter (when sending traffic to multiple destinations), since this algorithm uses multiple system-wide counters (rather than a single system-wide counter). The extent to which the re-use frequency will be lower will depend on the number of elements in counter[], and the number of other active flows that result in the same value of G() (and hence cause the same counter to be incremented for each fragmented datagram that is sent).
- o It is possible to implement the algorithm such that good performance is achieved. For example, the result of F() could be stored in the Destinations Cache (such that it need not be recomputed for each packet that must be sent) along with the computed "index"/argument for counter[].

NOTE: If this implementation approach is followed, and an entry of the Destinations Cache must be removed as a result of resource management, the last Fragment Identification value used for that Destination will **not** be lost. This is an improvement over the algorithm specified in [Section 6.1](#).

The possible drawbacks of this algorithm are:

- o Since the Fragment Identification values are predictable by the destination host, a vulnerable host could possibly leak to third-parties the Fragment Identification values used by other hosts to send traffic to it (i.e., Host B could leak to Host C the Fragment Identification values that Host A is using to send packets to Host B). [Appendix A](#) describes a possible scenario in which that information leakage could take place. We note, however, that this algorithm makes the aforementioned attack less reliable for the attacker, since each counter could be possibly shared by multiple traffic flows (i.e., packets destined to other destinations might cause the same counter to be incremented).

This algorithm might be preferable (over the one specified in [Section 6.1](#)) in those scenarios in which a node is expected to communicate with a large number of destinations, and thus it is desirable to limit the amount of information to be maintained in memory.

NOTE: In such scenarios, if the algorithm specified in [Section 6.1](#) were implemented, entries from the Destinations Cache might need to be pruned frequently, thus increasing the risk of fragment Identification collisions.

[7.](#) 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.

[8.](#) Security Considerations

This document discusses the security implications of predictable Fragment Identification values, and updates [RFC 2460](#) such that Fragment Identification values are required to be unpredictable by off-path attackers, hence mitigating the aforementioned security implications.

A number of possible algorithms are specified, to provide some implementation alternatives to implementers. However, the selection of a specific algorithm is left to implementers. We note that the selection of such an algorithm usually implies a number of trade-offs (security, performance, implementation complexity, interoperability properties, etc.).

[9.](#) Acknowledgements

The author would like to thank Ivan Arce for proposing the attack scenario described in [Appendix A](#).

The author would like to thank Ivan Arce and Dave Thaler for providing valuable comments on earlier versions of this document.

This document is based on the technical report "Security Assessment of the Internet Protocol version 6 (IPv6)" [[CPNI-IPv6](#)] authored by Fernando Gont on behalf of the UK Centre for the Protection of National Infrastructure (CPNI).

[10.](#) References

[10.1.](#) 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.
- [RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", [BCP 106](#), [RFC 4086](#), June 2005.
- [RFC5722] Krishnan, S., "Handling of Overlapping IPv6 Fragments", [RFC 5722](#), December 2009.
- [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-Protocol Port Randomization", [BCP 156](#), [RFC 6056](#), January 2011.
- [RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", [RFC 6145](#), April 2011.
- [RFC6946] Gont, F., "Processing of IPv6 "Atomic" Fragments", [RFC 6946](#), May 2013.

10.2. Informative References

- [RFC4963] Heffner, J., Mathis, M., and B. Chandler, "IPv4 Reassembly Errors at High Data Rates", [RFC 4963](#), July 2007.
- [RFC5927] Gont, F., "ICMP Attacks against TCP", [RFC 5927](#), July 2010.
- [RFC6274] Gont, F., "Security Assessment of the Internet Protocol Version 4", [RFC 6274](#), July 2011.
- [Bellovin2002] Bellovin, S., "A Technique for Counting NATted Hosts", IMW'02 Nov. 6-8, 2002, Marseille, France, 2002.
- [CPNI-IPv6] Gont, F., "Security Assessment of the Internet Protocol version 6 (IPv6)", UK Centre for the Protection of National Infrastructure, (available on request).
- [Fyodor2004] Fyodor, , "Idle scanning and related IP ID games", 2004, <<http://www.insecure.org/nmap/idlescan.html>>.
- [Herzberg2013] Herzberg, A. and H. Shulman, "Fragmentation Considered Poisonous", Technical Report 13-03, March 2013, <<http://u.cs.biu.ac.il/~herzbea/security/13-03-frag.pdf>>.

[Klein2007]

Klein, A., "OpenBSD DNS Cache Poisoning and Multiple O/S Predictable IP ID Vulnerability", 2007,
<http://www.trusteer.com/files/OpenBSD_DNS_Cache_Poisoning_and_Multiple_OS_Predictable_IP_ID_Vulnerability.pdf>.

[Sanfilippo1998a]

Sanfilippo, S., "about the ip header id", Post to Bugtraq mailing-list, Mon Dec 14 1998,
<<http://www.kyuzz.org/antirez/papers/ipid.html>>.

[Sanfilippo1998b]

Sanfilippo, S., "Idle scan", Post to Bugtraq mailing-list, 1998, <<http://www.kyuzz.org/antirez/papers/dumbscan.html>>.

[Sanfilippo1999]

Sanfilippo, S., "more ip id", Post to Bugtraq mailing-list, 1999,
<<http://www.kyuzz.org/antirez/papers/moreipid.html>>.

[SI6-IPv6]

"SI6 Networks' IPv6 toolkit",
<<http://www.si6networks.com/tools/ipv6toolkit>>.

[Zalewski2003]

Zalewski, M., "A new TCP/IP blind data injection technique?", Post to Bugtraq mailing-list, Thu, 11 Dec 2003 00:28:28 +0100 (CET), 2003,
<<http://lcamtuf.coredump.cx/ipfrag.txt>>.

Appendix A. Information leakage produced by vulnerable implementations

Section 3 provides a number of references describing a number of ways in which a vulnerable implementation may reveal the Fragment Identification values to be used in subsequent packets, thus opening the door to a number of attacks. In all of those scenarios, a vulnerable implementation leaks/reveals its own Identification number.

This section presents a different case, in which a vulnerable implementation leaks/reveals the Identification number of a non-vulnerable implementation. That is, a vulnerable implementation (Host A) leaks the current Fragment Identification value in use by a third-party host (Host B) to send fragmented datagrams from Host B to Host A.

For the most part, this section is included to illustrate how a vulnerable implementation might be leveraged to leak-out the

Fragment Identification value of an otherwise non-vulnerable implementation. This section might be removed in future revisions of this document.

The following scenarios assume:

Host A:

Is an IPv6 host that implements the recommended Fragment Identification algorithm ([Section 6.1](#)), implements [[RFC5722](#)], but does not implement [[RFC6946](#)].

Host B:

Victim node. Selected the Fragment Identification values from a global counter.

Host C:

Attacker. Can forge the IPv6 Source Address of his packets at will.

In the following scenarios, large ICMPv6 Echo Request packets are employed to "sample" the Fragment Identification value of a host. We note that while the figures show only one packet for the ICMPv6 Echo Request and the ICMPv6 Echo Response, each of those packets will typically comprise two fragments, such that the resulting datagram is larger than the MTU of the networks to which Host B and Host C are attached.

In the lines #1-#2 (and lines #8-#9), the attacker samples the current Fragment Identification value. In line #3, the attacker sends a forged TCP SYN segment to Host A. If corresponding TCP port is closed, and the attacker fails when trying to produce a collision of Fragment Identifications (see line #4), the following packet exchange might take place:

A	B	C
#1	<----- Echo Req #1 -----	
#2	--- Echo Resp #1, FID=5000 --->	
#3	<----- SYN #1, src= B -----	
#4	<--- SYN/ACK, FID=42 src = A---	
#5	---- SYN/ACK, FID=9000 ---->	
#6	<----- RST, FID= 5001 -----	
#7	<----- RST, FID= 5002 -----	
#8	<----- Echo Req #2 -----	
#9	--- Echo Resp #2, FID=5003 --->	

On the other hand, if the attacker succeeds to produce a collision of Fragment Identification values, the following packet exchange could take place:

```

      A                               B                               C
#1                                     <----- Echo Req #1 ----->
#2                                     --- Echo Resp #1, FID=5000 --->
#3 <----- SYN #1, src= B ----->
#4 <-- SYN/ACK, FID=9000 src=A ---
#5 ---- SYN/ACK, FID=9000 ---->
      ... (RFC5722) ...
#6                                     <----- Echo Req #2 ----->
#7                                     ---- Echo Resp #2, FID=5001 -->

```

Clearly, the Fragment Identification value sampled by from the second ICMPv6 Echo Response packet ("Echo Resp #2") implicitly indicates whether the Fragment Identification in the forged SYN/ACK (see line #4 in both figures) was the current Fragment Identification in use by Host A.

As a result, the attacker could employ this technique to learn the current Fragment Identification value used by host A to send packets to host B, even when Host A itself has a non-vulnerable implementation.

[Appendix B](#). Survey of Fragment Identification selection algorithms employed by popular IPv6 implementations

This section includes a survey of the Fragment Identification selection algorithms employed in some popular operating systems.

The survey was produced with the SI6 Networks IPv6 toolkit [[SI6-IPv6](#)].

Operating System	Algorithm
FreeBSD 9.0	Unpredictable (Random)
Linux 3.0.0-15	Predictable (Global Counter, Init=0, Incr=1)
Linux-current	Unpredictable (Per-dest Counter, Init=random, Incr=1)
NetBSD 5.1	Unpredictable (Random)
OpenBSD-current	Random (SKIP32)
Solaris 10	Predictable (Per-dst Counter, Init=0, Incr=1)
Windows XP SP2	Predictable (Global Counter, Init=0, Incr=2)
Windows Vista (Build 6000)	Predictable (Global Counter, Init=0, Incr=2)
Windows 7 Home Premium	Predictable (Global Counter, Init=0, Incr=2)

Table 1: Fragment Identification algorithms employed by different OSes

In the text above, "predictable" should be taken as "easily guessable by an off-path attacker, by sending a few probe packets".

Author's Address

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>