

Internet Research Task Force (IRTF)
Internet-Draft
Intended status: Informational
Expires: July 17, 2021

F. Gont
SI6 Networks
I. Arce
Quarkslab
January 13, 2021

**Unfortunate History of Transient Numeric Identifiers
draft-irtf-pearg-numeric-ids-history-06**

Abstract

This document analyzes the timeline of the specification and implementation of different types of "transient numeric identifiers" used in IETF protocols, and how the security and privacy properties of such protocols have been affected as a result of it. It provides empirical evidence that advice in this area is warranted. This document is a product of the Privacy Enhancement and Assessment Research Group (PEARG) in the IRTF.

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 <https://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 July 17, 2021.

Copyright Notice

Copyright (c) 2021 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 (<https://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.

Table of Contents

1.	Introduction	2
2.	Terminology	4
3.	Threat Model	4
4.	Issues with the Specification of Transient Numeric Identifiers	4
5.	IPv4/IPv6 Identification	5
6.	TCP Initial Sequence Numbers (ISNs)	9
7.	IPv6 Interface Identifiers (IIDs)	11
8.	NTP Reference IDs (REFIDs)	14
9.	Transport Protocol Ephemeral Port Numbers	15
10.	IANA Considerations	16
11.	Security Considerations	16
12.	Acknowledgements	17
13.	References	17
13.1.	Normative References	17
13.2.	Informative References	19
	Authors' Addresses	28

[1.](#) Introduction

Networking protocols employ a variety of transient numeric identifiers for different protocol objects, such as IPv4 and IPv6 Fragment Identifiers [[RFC0791](#)] [[RFC8200](#)], IPv6 Interface Identifiers (IIDs) [[RFC4291](#)], transport protocol ephemeral port numbers [[RFC6056](#)], TCP Initial Sequence Numbers (ISNs) [[RFC0793](#)], and DNS Transaction IDs (TxIDs) [[RFC1035](#)]. These identifiers typically have specific interoperability requirements (e.g. uniqueness during a specified period of time), and associated failure modes when such requirements are not met [[I-D.irtf-pearg-numeric-ids-generation](#)].

For more than 30 years, a large number of implementations of the TCP/IP protocol suite have been subject to a variety of attacks, with effects ranging from Denial of Service (DoS) or data injection, to information leakages that could be exploited for pervasive monitoring [[RFC7258](#)]. The root cause of these issues has been, in many cases, poor selection of transient numeric identifiers, usually as a result of insufficient or misleading specifications.

For example, implementations have been subject to security or privacy issues resulting from:

- o Predictable IPv4 or IPv6 Fragment Identifiers (see e.g. [[Sanfilippo1998a](#)], [[RFC6274](#)], and [[RFC7739](#)])
- o Predictable IPv6 IIDs (see e.g. [[RFC7721](#)], [[RFC7707](#)], and [[RFC7217](#)])

- o Predictable transport protocol ephemeral port numbers (see e.g. [[RFC6056](#)] and [[Silbersack2005](#)])
- o Predictable TCP Initial Sequence Numbers (ISNs) (see e.g. [[Morris1985](#)], [[Bellovin1989](#)], and [[RFC6528](#)])
- o Predictable DNS TxIDs (see e.g. [[Schuba1993](#)] and [[Klein2007](#)])

These examples indicate that when new protocols are standardized or implemented, the security and privacy properties of the associated transient numeric identifiers tend to be overlooked, and inappropriate algorithms to generate such identifiers (i.e. that negatively affect the security or privacy properties of the protocol) are either suggested in the specification or selected by implementers.

This document contains a non-exhaustive timeline of the specification and vulnerability disclosures related to some sample transient numeric identifiers, including other work that has led to advances in this area. This analysis indicates that:

- o Vulnerabilities associated with the inappropriate generation of transient numeric identifiers have affected protocol implementations for an extremely long period of time.
- o Such vulnerabilities, even when addressed for a given protocol version, were later reintroduced in new versions or new implementations of the same protocol.
- o Standardization efforts that discuss and provide advice in this area can have a positive effect on protocol specifications and protocol implementations.

While it is generally possible to identify an algorithm that can satisfy the interoperability requirements for a given transient numeric identifier, this document provides empirical evidence that doing so without negatively affecting the security or privacy properties of the aforementioned protocols is non-trivial. Other related documents ([[I-D.irtf-pearg-numeric-ids-generation](#)] and [[I-D.gont-numeric-ids-sec-considerations](#)]) provide guidance in this area, as motivated by the present document.

This document represents the consensus of the Privacy Enhancement and Assessment Research Group (PEARG).

2. Terminology

Transient Numeric Identifier:

A data object in a protocol specification that can be used to definitely distinguish a protocol object (a datagram, network interface, transport protocol endpoint, session, etc) from all other objects of the same type, in a given context. Transient numeric identifiers are usually defined as a series of bits, and represented using integer values. These identifiers are typically dynamically selected, as opposed to statically-assigned numeric identifiers (see e.g. [[IANA-PROT](#)]). We note that different identifiers may have additional requirements or properties depending on their specific use in a protocol. We use the term "transient numeric identifier" (or simply "numeric identifier" or "identifier" as short forms) as a generic term to refer to any data object in a protocol specification that satisfies the identification property stated above.

The terms "constant IID", "stable IID", and "temporary IID" are to be interpreted as defined in [[RFC7721](#)].

3. Threat Model

Throughout this document, we assume an attacker does not have physical or logical access to the system(s) being attacked, and cannot observe the packets being transferred between the sender and the receiver(s) of the target protocol (if any). However, we assume the attacker can send any traffic to the target device(s), to e.g. sample transient numeric identifiers employed by such device(s).

4. Issues with the Specification of Transient Numeric Identifiers

While assessing protocol specifications regarding the use of transient numeric identifiers, we have found that most of the issues discussed in this document arise as a result of one of the following conditions:

- o Protocol specifications that under-specify the requirements for their transient numeric identifiers
- o Protocol specifications that over-specify their transient numeric identifiers
- o Protocol implementations that simply fail to comply with the specified requirements

A number of protocol specifications (too many of them) have simply overlooked the security and privacy implications of transient numeric

identifiers. Examples of them are the specification of TCP ephemeral ports in [\[RFC0793\]](#), the specification of TCP sequence numbers in [\[RFC0793\]](#), or the specification of the DNS TxID in [\[RFC1035\]](#).

On the other hand, there are a number of protocol specifications that over-specify some of their associated transient numeric identifiers. For example, [\[RFC4291\]](#) essentially overloads the semantics of IPv6 Interface Identifiers (IIDs) by embedding link-layer addresses in the IPv6 IIDs, when the interoperability requirement of uniqueness could be achieved in other ways that do not result in negative security and privacy implications [\[RFC7721\]](#). Similarly, [\[RFC2460\]](#) suggested the use of a global counter for the generation of Fragment Identification values, when the interoperability properties of uniqueness per {Src IP, Dst IP} could be achieved with other algorithms that do not result in negative security and privacy implications [\[RFC7739\]](#).

Finally, there are protocol implementations that simply fail to comply with existing protocol specifications. For example, some popular operating systems (notably Microsoft Windows) still fail to implement transport protocol ephemeral port randomization, as recommended in [\[RFC6056\]](#).

5. IPv4/IPv6 Identification

This section presents the timeline of the Identification field employed by IPv4 (in the base header) and IPv6 (in Fragment Headers). The reason for presenting both cases in the same section is to make it evident that while the Identification value serves the same purpose in both IPv4 and IPv6, the work and research done for the IPv4 case did not affect IPv6 specifications or implementations.

The IPv4 Identification is specified in [\[RFC0791\]](#), which specifies the interoperability requirements for the Identification field: the sender must choose the Identification field to be unique for a given source address, destination address, and protocol, for the time the datagram (or any fragment of it) could be alive in the internet. It suggests that a node may keep "a table of Identifiers, one entry for each destination it has communicated with in the last maximum packet lifetime for the internet", and suggests that "since the Identifier field allows 65,536 different values, hosts may be able to simply use unique identifiers independent of destination". The above has been interpreted numerous times as a suggestion to employ per-destination or global counters for the generation of Identification values. While [\[RFC0791\]](#) does not suggest any flawed algorithm for the generation of Identification values, the specification omits a discussion of the security and privacy implications of predictable Identification values. This has resulted in many IPv4

implementations generating predictable fragment Identification values by means of a global counter, at least at some point in time.

The IPv6 Identification was originally specified in [\[RFC1883\]](#). It serves the same purpose as its IPv4 counterpart, with the only difference residing in the length of the corresponding field, and that while the IPv4 Identification field is part of the base IPv4 header, in the IPv6 case it is part of the Fragment header (which may or may not be present in an IPv6 packet). [\[RFC1883\]](#) states, in [Section 4.5](#), that the Identification must be different than that of any other fragmented packet sent recently (within the maximum likely lifetime of a packet) with the same Source Address and Destination Address. Subsequently, it notes that this requirement can be met by means of a wrap-around 32-bit counter that is incremented each time a packet must be fragmented, and that it is an implementation choice whether to use a global or a per-destination counter. Thus, the implementation of the IPv6 Identification is similar to that of the IPv4 case, with the only difference that in the IPv6 case the suggestions to use simple counters is more explicit. [\[RFC2460\]](#) was the first revision of the core IPv6 specification, and maintained the same text for the specification of the IPv6 Identification field. [\[RFC8200\]](#), the second revision of the core IPv6 specification, removes the suggestion from [\[RFC2460\]](#) to use a counter for the generation of IPv6 Identification values, and points to [\[RFC7739\]](#) for sample algorithms for their generation.

September 1981:

[\[RFC0791\]](#) specifies the interoperability requirements for IPv4 Identification value, but does not perform a vulnerability assessment of this transient numeric identifier.

December 1995:

[\[RFC1883\]](#), the first specification of the IPv6 protocol, is published. It suggests that a counter be used to generate the IPv6 Identification value, and notes that it is an implementation choice whether to maintain a single counter for the node or multiple counters, e.g., one for each of the node's possible source addresses, or one for each active (source address, destination address) combination.

December 1998:

[\[Sanfilippo1998a\]](#) finds that predictable IPv4 Identification values (generated by most popular implementations) can be leveraged to count the number of packets sent by a target node. [\[Sanfilippo1998b\]](#) explains how to leverage the same vulnerability to implement a port-scanning technique known as dumb/idle scan. A tool that implements this attack is publicly released.

December 1998:

[[RFC2460](#)], a revision of the IPv6 specification, is published, obsoleting [[RFC1883](#)]. It maintains the same specification of the IPv6 Identification field as its predecessor ([[RFC1883](#)]).

November 1999:

[[Sanfilippo1999](#)] discusses how to leverage predictable IPv4 Identification to uncover the rules of a number of firewalls.

December 1998:

OpenBSD implements randomization of the IPv4 Identification field [[OpenBSD-IPv4-ID](#)]. This feature eventually shipped with OpenBSD 2.5.

November 1999:

[[Bellovin2002](#)] explains how the IPv4 Identification field can be exploited to count the number of systems behind a NAT.

September 2002:

[[Fyodor2002](#)] explains how to implement a stealth port-scanning technique by leveraging nodes that employ predictable IPv4 Identification values.

October 2003:

OpenBSD implements randomization of the IPv6 Identification field [[OpenBSD-IPv6-ID](#)]. This feature eventually shipped with OpenBSD 3.4.

December 2003:

[[Zalewski2003](#)] explains a technique to perform TCP data injection attack based on predictable IPv4 identification values which requires less effort than TCP injection attacks performed with bare TCP packets.

November 2005:

[[Silbersack2005](#)] discusses shortcoming in a number of techniques to mitigate predictable IPv4 Identification values.

October 2007:

[[Klein2007](#)] describes a weakness in the pseudo random number generator (PRNG) in use for the generation of the IP Identification by a number of operating systems.

June 2011:

[[Gont2011](#)] describes how to perform idle scan attacks in IPv6.

November 2011:

Linux mitigates predictable IPv6 Identification values
[[RedHat2011](#)] [[SUSE2011](#)] [[Ubuntu2011](#)].

December 2011:

[[draft-gont-6man-predictable-fragment-id-00](#)] describes the security implications of predictable IPv6 Identification values, and possible mitigations. This document has the Intended Status of "Standards Track", with the intention to formally update [[RFC2460](#)], to introduce security and privacy requirements on IPv6 Identification values.

May 2012:

[[Gont2012](#)] notes that some major IPv6 implementations still employ predictable IPv6 Identification values.

March 2013:

The 6man WG adopts [[I-D.gont-6man-predictable-fragment-id](#)], but changes the track to "BCP" (while still formally updating [[RFC2460](#)]), publishing the resulting document as [[draft-ietf-6man-predictable-fragment-id-00](#)].

June 2013:

A patch that implements IPv6-based idle-scan in nmap is submitted [[Morbitzer2013](#)].

December 2014:

The 6man WG changes the Intended Status of [[draft-ietf-6man-predictable-fragment-id-01](#)] to "Informational" and publishes it as [[draft-ietf-6man-predictable-fragment-id-02](#)]. As a result, it no longer formally updates [[RFC2460](#)].

June 2015:

[[draft-ietf-6man-predictable-fragment-id-08](#)] notes that some popular host and router implementations still employ predictable IPv6 Identification values.

February 2016:

[[RFC7739](#)] (based on [[I-D.ietf-6man-predictable-fragment-id](#)]) analyzes the security and privacy implications of predictable IPv6 Identification values, and provides guidance for selecting an algorithm to generate such values. However, being published with the Intended Status of "Informational", it does not formally update [[RFC2460](#)].

June 2016:

[[I-D.ietf-6man-rfc2460bis](#)], revision of [[RFC2460](#)], removes the suggestion from [RFC2460](#) to use a counter for the generation of

IPv6 Identification values, but does not perform a vulnerability assessment of the generation of IPv6 Identification values.

July 2017:

[[I-D.ietf-6man-rfc2460bis](#)] is finally published as [[RFC8200](#)], obsoleting [[RFC2460](#)], and pointing to [[RFC7739](#)] for sample algorithms for the generation of IPv6 Fragment Identification values.

June 2019:

[[IPID-DEV](#)] notes that the IPv6 ID generators of two popular operating systems are flawed.

6. TCP Initial Sequence Numbers (ISNs)

[RFC0793] suggests that the choice of the ISN of a connection is not arbitrary, but aims to reduce the chances of a stale segment from being accepted by a new incarnation of a previous connection.

[RFC0793] suggests the use of a global 32-bit ISN generator that is incremented by 1 roughly every 4 microseconds. However, as a matter of fact, protection against stale segments from a previous incarnation of the connection is enforced by preventing the creation of a new incarnation of a previous connection before $2 \times \text{MSL}$ have passed since a segment corresponding to the old incarnation was last seen (where "MSL" is the "Maximum Segment Lifetime" [[RFC0793](#)]). This is accomplished by the TIME-WAIT state and TCP's "quiet time" concept (see [Appendix B of \[RFC1323\]](#)). Based on the assumption that ISNs are monotonically increasing across connections, many stacks (e.g., 4.2BSD-derived) use the ISN of an incoming SYN segment to perform "heuristics" that enable the creation of a new incarnation of a connection while the previous incarnation is still in the TIME-WAIT state (see p. 945 of [[Wright1994](#)]). This avoids an interoperability problem that may arise when a node establishes connections to a specific TCP end-point at a high rate [[Silbersack2005](#)].

The interoperability requirements for TCP ISNs are probably not as clearly spelled out as one would expect. Furthermore, the suggestion of employing a global counter in [[RFC0793](#)] negatively affects the security and privacy properties of the protocol.

September 1981:

[[RFC0793](#)], suggests the use of a global 32-bit ISN generator, whose lower bit is incremented roughly every 4 microseconds. However, such an ISN generator makes it trivial to predict the ISN that a TCP instance will use for new connections, thus allowing a variety of attacks against TCP.

February 1985:

[Morris1985] was the first to describe how to exploit predictable TCP ISNs for forging TCP connections that could then be leveraged for trust relationship exploitation.

April 1989:

[Bellovin1989] discussed the security considerations for predictable ISNs (along with a range of other protocol-based vulnerabilities).

February 1995:

[Shimomura1995] reported a real-world exploitation of the attack described in 1985 (ten years before) in [Morris1985].

May 1996:

[RFC1948] was the first IETF effort, authored by Steven Bellovin, to address predictable TCP ISNs. The same concept specified in this document for TCP ISNs was later proposed for TCP ephemeral ports [RFC6056], TCP Timestamps, and eventually even IPv6 Interface Identifiers [RFC7217].

July 1996:

OpenBSD implements TCP ISN randomization based on random increments (please see [Appendix A.2](#) of [I-D.irtf-pearg-numeric-ids-generation]) [OpenBSD-TCP-ISN-I]. The feature eventually shipped with OpenBSD 2.0.

December 2000:

OpenBSD implements TCP ISN randomization using simple randomization (please see Section 7.1 of [I-D.irtf-pearg-numeric-ids-generation]) [OpenBSD-TCP-ISN-R]. The feature eventually shipped with OpenBSD 2.9.

March 2001:

[Zalewski2001] provides a detailed analysis of statistical weaknesses in some ISN generators, and includes a survey of the algorithms in use by popular TCP implementations.

May 2001:

Vulnerability advisories [CERT2001] [USCERT2001] are released regarding statistical weaknesses in some ISN generators, affecting popular TCP/IP implementations.

March 2002:

[Zalewski2002] updates and complements [Zalewski2001]. It concludes that "while some vendors [...] reacted promptly and tested their solutions properly, many still either ignored the issue and never evaluated their implementations, or implemented a

flawed solution that apparently was not tested using a known approach" [[Zalewski2002](#)].

June 2007:

OpenBSD implements TCP ISN randomization based on the algorithm specified in [[RFC1948](#)] (currently obsoleted by [[RFC6528](#)]) for the TCP endpoint that performs the active open, while keeping the simple randomization scheme for the endpoint performing the passive open [[OpenBSD-TCP-ISN-H](#)]. This provides monotonically-increasing ISNs for the client side (allowing the BSD heuristics to work as expected), while avoiding any patterns in the ISN generation for the server side. This feature eventually shipped with OpenBSD 4.2.

February 2012:

[[RFC6528](#)], published 27 years after Morris' original work [[Morris1985](#)], formally updates [[RFC0793](#)] to mitigate predictable TCP ISNs.

August 2014:

[[I-D.eddy-rfc793bis-04](#)], the upcoming revision of the core TCP protocol specification, incorporates the algorithm specified in [[RFC6528](#)] as the recommended algorithm for TCP ISN generation.

7. IPv6 Interface Identifiers (IIDs)

IPv6 Interface Identifiers can be generated in multiple ways: SLAAC [[RFC4862](#)], DHCPv6 [[RFC8415](#)], and manual configuration. This section focuses on Interface Identifiers resulting from SLAAC.

The Interface Identifier of stable (traditional) IPv6 addresses resulting from SLAAC have traditionally resulted in the underlying link-layer address being embedded in the IID. At the time, employing the underlying link-layer address for the IID was seen as a convenient way to obtain a unique address. However, recent awareness about the security and privacy properties of this approach [[RFC7707](#)] [[RFC7721](#)] has led to the replacement of this flawed scheme with an alternative one [[RFC7217](#)] [[RFC8064](#)] that does not negatively affect the security and privacy properties of the protocol.

January 1997:

[[RFC2073](#)] specifies the syntax of IPv6 global addresses (referred to as "An IPv6 Provider-Based Unicast Address Format" at the time), consistent with the IPv6 addressing architecture specified in [[RFC1884](#)]. Hosts are recommended to "generate addresses using link-specific addresses as Interface ID such as 48 bit IEEE-802 MAC addresses".

July 1998:

[[RFC2374](#)] specifies "An IPv6 Aggregatable Global Unicast Address Format" (obsoleting [[RFC2373](#)]) changing the size of the Interface ID to 64 bits, and specifies that that IIDs must be constructed in IEEE EUI-64 format. How such identifiers are constructed becomes specified in the appropriate "IPv6 over <link>" specification such as "IPv6 over Ethernet".

January 2001:

[[RFC3041](#)] recognizes the problem of network activity correlation, and specifies temporary addresses. Temporary addresses are to be used along with stable addresses.

August 2003:

[[RFC3587](#)] obsoletes [[RFC2374](#)], making the TLA/NLA structure historic. The syntax and recommendations for the traditional stable IIDs remain unchanged, though.

February 2006:

[[RFC4291](#)] is published as the latest "IP Version 6 Addressing Architecture", requiring the IIDs of the traditional (stable) autoconfigured addresses to employ the Modified EUI-64 format. The details of constructing such interface identifiers are defined in the appropriate "IPv6 over <link>" specifications.

March 2008:

[[RFC5157](#)] provides hints regarding how patterns in IPv6 addresses could be leveraged for the purpose of address scanning.

December 2011:

[[draft-gont-6man-stable-privacy-addresses-00](#)] notes that the traditional scheme for generating stable addresses allows for address scanning, and also does not prevent active node tracking. It also specifies an alternative algorithm meant to replace IIDs based on Modified EUI-64 format identifiers.

November 2012:

The 6man WG adopts [[I-D.gont-6man-stable-privacy-addresses](#)] as a working group item (as [[draft-ietf-6man-stable-privacy-addresses-00](#)]). However, the document no longer formally updates [[RFC4291](#)], and therefore the specified algorithm no longer formally replaces the Modified EUI-64 format identifiers.

February 2013:

An address-scanning tool (scan6 of [[IPv6-Toolkit](#)]) that leverages IPv6 address patterns is released [[Gont2013](#)].

July 2013:

[[I-D.cooper-6man-ipv6-address-generation-privacy](#)] elaborates on the security and privacy properties of all known algorithms for generating IPv6 IIDs.

January 2014:

The 6man WG publishes [[draft-ietf-6man-default-iids-00](#)] ("Recommendation on Stable IPv6 Interface Identifiers"), recommending [[I-D.ietf-6man-stable-privacy-addresses](#)] for the generation of stable addresses.

April 2014:

[[RFC7217](#)] (formerly [[I-D.ietf-6man-stable-privacy-addresses](#)]) is published, specifying "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)" as an alternative to (but *not* replacement of) Modified EUI-64 format IIDs.

March 2016:

[[RFC7707](#)] (formerly [[I-D.gont-opsec-ipv6-host-scanning](#)], and later [[I-D.ietf-opsec-ipv6-host-scanning](#)]), about "Network Reconnaissance in IPv6 Networks", is published.

March 2016:

[[RFC7721](#)] (formerly [[I-D.cooper-6man-ipv6-address-generation-privacy](#)] and later [[I-D.ietf-6man-ipv6-address-generation-privacy](#)]), about "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", is published.

May 2016:

[[draft-gont-6man-non-stable-iids-00](#)] is published, with the goal of specifying requirements for non-stable addresses, and updating [[RFC4941](#)] such that use of only temporary addresses is allowed.

May 2016:

[[draft-gont-6man-address-usage-recommendations-00](#)] is published, providing an analysis of how different aspects on an address (from stability to usage mode) affect their corresponding security and privacy properties, and meaning to eventually provide advice in this area.

February 2017:

The 6man WG publishes [[RFC8064](#)] ("Recommendation on Stable IPv6 Interface Identifiers") (formerly [[I-D.ietf-6man-default-iids](#)]), with requirements for stable addresses and a recommendation to employ [[RFC7217](#)] for the generation of stable addresses. It formally updates a large number of RFCs.

March 2018:

[[draft-fgont-6man-rfc4941bis-00](#)] is published (as suggested by the 6man WG), to address flaws in [[RFC4941](#)] by revising it (as an alternative to the [[draft-gont-6man-non-stable-iids-00](#)] effort, published in March 2016).

July 2018:

[[draft-ietf-6man-rfc4941bis-00](#)] is adopted (as [[draft-fgont-6man-rfc4941bis-00](#)]) as a WG item of the 6man WG.

March 2020:

[[I-D.ietf-6man-rfc4941bis](#)] passes WGLC.

December 2020:

[[I-D.ietf-6man-rfc4941bis](#)] is finally approved by the IESG for publication as an RFC.

8. NTP Reference IDs (REFIDs)

NTP [[RFC5905](#)] Reference IDs are employed to avoid degree-one timing loops in scenarios where two NTP peers are (mutually) the time source of each other.

June 2010:

[[RFC5905](#)], "Network Time Protocol Version 4: Protocol and Algorithms Specification" is published. It specifies that for NTP peers with stratum higher than 1 the REFID embeds the IPv4 Address of the time source or an MD5 hash of the IPv6 address of the time source.

July 2016:

[[draft-stenn-ntp-not-you-refid-00](#)] is published, describing the information leakage produced via the NTP REFID. It proposes that NTP returns a special REFID when a packet employs an IP Source Address that is not believed to be a current NTP peer, but otherwise generates and returns the traditional REFID. It is subsequently adopted by the NTP WG as [[I-D.ietf-ntp-refid-updates](#)].

April 2019:

[[Gont-NTP](#)] notes that the proposed fix specified in [[draft-ietf-ntp-refid-updates-00](#)] is, at the very least, sub-optimal.

9. Transport Protocol Ephemeral Port Numbers

Most (if not all) transport protocols employ "port numbers" to demultiplex packets to the corresponding transport protocol instances.

August 1980:

[[RFC0768](#)] notes that the UDP source port is optional and identifies the port of the sending process. It does not specify interoperability requirements for source port selection, nor does it suggest possible ways to select port numbers. Most popular implementations end up selecting source ports from a system-wide global counter.

September 1981:

[[RFC0793](#)] (the TCP specification) essentially describes the use of port numbers, and specifies that port numbers should result in a unique socket pair (local address, local port, remote address, remote port). How ephemeral ports (i.e. port numbers for "active opens") are selected, and the port range from which they are selected, are left unspecified.

July 1996:

OpenBSD implements ephemeral port randomization [[OpenBSD-PR](#)]. The feature eventually shipped with OpenBSD 2.0.

July 2008:

The CERT Coordination Centre published details of what became known as the "Kaminsky Attack" [[VU-800113](#)] on the DNS. The attack exploited the lack of source port randomisation in many major DNS implementations to perform cache poisoning in an effective and practical manner.

January 2009:

[[RFC5452](#)] mandates the use of port randomization for DNS resolvers, and mandates that implementations must randomize ports from the range (53 or 1024, and above) or the largest possible port range. It does not recommend possible algorithms for port randomization, although the document specifically targets DNS resolvers, for which a simple port randomization suffices (e.g. Algorithm 1 of [[RFC6056](#)]). This document led to the implementation of port randomization in the DNS resolver themselves, rather than in the underlying transport-protocols.

January 2011:

[[RFC6056](#)] notes that many TCP and UDP implementations result in predictable port numbers, and also notes that many implementations select port numbers from a small portion of the whole port number

space. It recommends the implementation and use of ephemeral port randomization, proposes a number of possible algorithms for port randomization, and also recommends to randomize port numbers over the range 1024-65535.

March 2016:

[[NIST-NTP](#)] reports a non-normal distribution of the ephemeral port numbers employed by the NTP clients of an Internet Time Service.

April 2019:

[[I-D.gont-ntp-port-randomization](#)] notes that some NTP implementations employ the NTP service port (123) as the local port for non-symmetric modes, and aims to update the NTP specification to recommend port randomization in such cases, in line with [[RFC6056](#)]. The proposal experiences some push-back in the relevant working group (NTP WG) [[NTP-PORTR](#)], but is finally adopted as a working group item as [[I-D.ietf-ntp-port-randomization](#)].

September 2019:

[[I-D.ietf-ntp-port-randomization](#)] passes its WGLC.

10. 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.

11. Security Considerations

This document analyzes the timeline of the specification and implementation of the transient numeric identifiers of some sample IETF protocols, and how the security and privacy properties of such protocols have been affected as a result of it. It provides concrete evidence that advice in this area is warranted.

[[I-D.gont-numeric-ids-sec-considerations](#)] formally requires protocol specifications to specify the interoperability requirements for their transient numeric identifiers, to do a warranted vulnerability assessment of such transient numeric identifiers, and to recommend possible algorithms for their generation, such that the interoperability requirements are complied with, while any negative security and privacy properties of these transient numeric identifiers are mitigated.

[[I-D.irtf-pearg-numeric-ids-generation](#)] analyzes and categorizes transient numeric identifiers based on their interoperability requirements and their associated failure modes, and recommends

possible algorithms that can comply with those requirements without negatively affecting the security and privacy properties of the corresponding protocols.

12. Acknowledgements

The authors would like to thank (in alphabetical order) Bernard Aboba, Dave Crocker, Theo de Raadt, Sara Dickinson, Guillermo Gont, Christian Huitema, Kris Shrishak, Joe Touch, and Christopher Wood, for providing valuable comments on earlier versions of this document.

The authors would like to thank (in alphabetical order) Steven Bellovin, Joseph Lorenzo Hall, Gre Norcie, and Martin Thomson, for providing valuable comments on [[I-D.gont-predictable-numeric-ids](#)], on which this document is based.

[Section 6](#) of this document borrows text from [[RFC6528](#)], authored by Fernando Gont and Steven Bellovin.

The authors would like to thank Sara Dickinson and Christopher Wood, for their guidance during the publication process of this document.

The authors would like to thank Diego Armando Maradona for his magic and inspiration.

13. References

13.1. Normative References

- [RFC0768] Postel, J., "User Datagram Protocol", STD 6, [RFC 768](#), DOI 10.17487/RFC0768, August 1980, <<https://www.rfc-editor.org/info/rfc768>>.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.
- [RFC1323] Jacobson, V., Braden, R., and D. Borman, "TCP Extensions for High Performance", [RFC 1323](#), DOI 10.17487/RFC1323, May 1992, <<https://www.rfc-editor.org/info/rfc1323>>.
- [RFC1883] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 1883](#), DOI 10.17487/RFC1883, December 1995, <<https://www.rfc-editor.org/info/rfc1883>>.

- [RFC1884] Hinden, R., Ed. and S. Deering, Ed., "IP Version 6 Addressing Architecture", [RFC 1884](#), DOI 10.17487/RFC1884, December 1995, <<https://www.rfc-editor.org/info/rfc1884>>.
- [RFC2073] Rekhter, Y., Lothberg, P., Hinden, R., Deering, S., and J. Postel, "An IPv6 Provider-Based Unicast Address Format", [RFC 2073](#), DOI 10.17487/RFC2073, January 1997, <<https://www.rfc-editor.org/info/rfc2073>>.
- [RFC2373] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 2373](#), DOI 10.17487/RFC2373, July 1998, <<https://www.rfc-editor.org/info/rfc2373>>.
- [RFC2374] Hinden, R., O'Dell, M., and S. Deering, "An IPv6 Aggregatable Global Unicast Address Format", [RFC 2374](#), DOI 10.17487/RFC2374, July 1998, <<https://www.rfc-editor.org/info/rfc2374>>.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), DOI 10.17487/RFC2460, December 1998, <<https://www.rfc-editor.org/info/rfc2460>>.
- [RFC3041] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", [RFC 3041](#), DOI 10.17487/RFC3041, January 2001, <<https://www.rfc-editor.org/info/rfc3041>>.
- [RFC3587] Hinden, R., Deering, S., and E. Nordmark, "IPv6 Global Unicast Address Format", [RFC 3587](#), DOI 10.17487/RFC3587, August 2003, <<https://www.rfc-editor.org/info/rfc3587>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", [RFC 4941](#), DOI 10.17487/RFC4941, September 2007, <<https://www.rfc-editor.org/info/rfc4941>>.

- [RFC5452] Hubert, A. and R. van Mook, "Measures for Making DNS More Resilient against Forged Answers", [RFC 5452](#), DOI 10.17487/RFC5452, January 2009, <<https://www.rfc-editor.org/info/rfc5452>>.
- [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-Protocol Port Randomization", [BCP 156](#), [RFC 6056](#), DOI 10.17487/RFC6056, January 2011, <<https://www.rfc-editor.org/info/rfc6056>>.
- [RFC6528] Gont, F. and S. Bellovin, "Defending against Sequence Number Attacks", [RFC 6528](#), DOI 10.17487/RFC6528, February 2012, <<https://www.rfc-editor.org/info/rfc6528>>.
- [RFC7217] Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", [RFC 7217](#), DOI 10.17487/RFC7217, April 2014, <<https://www.rfc-editor.org/info/rfc7217>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, [RFC 8200](#), DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.
- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 8415](#), DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.

13.2. Informative References

- [Bellovin1989] Bellovin, S., "Security Problems in the TCP/IP Protocol Suite", Computer Communications Review, vol. 19, no. 2, pp. 32-48, 1989, <<https://www.cs.columbia.edu/~smb/papers/ipext.pdf>>.
- [Bellovin2002] Bellovin, S., "A Technique for Counting NATted Hosts", IMW'02 Nov. 6-8, 2002, Marseille, France, 2002.
- [CERT2001] CERT, "CERT Advisory CA-2001-09: Statistical Weaknesses in TCP/IP Initial Sequence Numbers", 2001, <https://resources.sei.cmu.edu/asset_files/WhitePaper/2001_019_001_496192.pdf>.

[[draft-fgont-6man-rfc4941bis-00](#)]

Gont, F., Krishnan, S., Narten, T., and R. Draves,
"Privacy Extensions for Stateless Address
Autoconfiguration in IPv6", [draft-fgont-6man-rfc4941bis-00](#)
(work in progress), March 2018.

[[draft-gont-6man-address-usage-recommendations-00](#)]

Gont, F. and W. Liu, "IPv6 Address Usage Recommendations",
[draft-gont-6man-address-usage-recommendations-00](#) (work in
progress), May 2016.

[[draft-gont-6man-non-stable-iids-00](#)]

Gont, F. and W. Liu, "Recommendation on Non-Stable IPv6
Interface Identifiers", [draft-gont-6man-non-stable-iids-00](#)
(work in progress), May 2016.

[[draft-gont-6man-predictable-fragment-id-00](#)]

Gont, F., "Security Implications of Predictable Fragment
Identification Values", [draft-gont-6man-predictable-
fragment-id-00](#) (work in progress), December 2011.

[[draft-gont-6man-stable-privacy-addresses-00](#)]

Gont, F., "A method for Generating Stable Privacy-Enhanced
Addresses with IPv6 Stateless Address Autoconfiguration
(SLAAC)", [draft-gont-6man-stable-privacy-addresses-00](#)
(work in progress), December 2011.

[[draft-ietf-6man-default-iids-00](#)]

Gont, F., Cooper, A., Thaler, D., and W. Liu,
"Recommendation on Stable IPv6 Interface Identifiers",
[draft-ietf-6man-default-iids-00](#) (work in progress), July
2014.

[[draft-ietf-6man-predictable-fragment-id-00](#)]

Gont, F., "Security Implications of Predictable Fragment
Identification Values", [draft-ietf-6man-predictable-
fragment-id-00](#) (work in progress), March 2013.

[[draft-ietf-6man-predictable-fragment-id-01](#)]

Gont, F., "Security Implications of Predictable Fragment
Identification Values", [draft-ietf-6man-predictable-
fragment-id-01](#) (work in progress), April 2014.

[[draft-ietf-6man-predictable-fragment-id-02](#)]

Gont, F., "Security Implications of Predictable Fragment
Identification Values", [draft-ietf-6man-predictable-
fragment-id-02](#) (work in progress), December 2014.

[[draft-ietf-6man-predictable-fragment-id-08](#)]

Gont, F., "Security Implications of Predictable Fragment Identification Values", [draft-ietf-6man-predictable-fragment-id-08](#) (work in progress), June 2015.

[[draft-ietf-6man-rfc4941bis-00](#)]

Gont, F., Krishnan, S., Narten, T., and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", [draft-ietf-6man-rfc4941bis-00](#) (work in progress), July 2018.

[[draft-ietf-6man-stable-privacy-addresses-00](#)]

Gont, F., "A method for Generating Stable Privacy-Enhanced Addresses with IPv6 Stateless Address Autoconfiguration (SLAAC)", [draft-ietf-6man-stable-privacy-addresses-00](#) (work in progress), May 2012.

[[draft-ietf-ntp-refid-updates-00](#)]

Goldberg, S. and H. Stenn, "Network Time Protocol Not You REFID", [draft-ietf-ntp-refid-updates-00](#) (work in progress), November 2016.

[[draft-stenn-ntp-not-you-refid-00](#)]

Goldberg, S. and S. KrishnanSTENN, "Network Time Protocol Not You REFID", [draft-stenn-ntp-not-you-refid-00](#) (work in progress), July 2016.

[Fyodor2002]

Fyodor, "Idle scanning and related IP ID games", 2002, <<http://www.insecure.org/nmap/idslescan.html>>.

[Gont-NTP]

Gont, F., "[Ntp] Comments on [draft-ietf-ntp-refid-updates-05](#)", Post to the NTP WG mailing list Message-ID: <d871d66d-4043-d8d0-f924-2191ebb2e2ce@si6networks.com>, April 2019, <<https://mailarchive.ietf.org/arch/msg/ntp/NkfTHxUU0dp14Agh3h1IPqfcRRg>>.

[Gont2011]

Gont, F., "Hacking IPv6 Networks (training course)", Hack In Paris 2011 Conference Paris, France, June 2011.

[Gont2012]

Gont, F., "Recent Advances in IPv6 Security", BSDCan 2012 Conference Ottawa, Canada. May 11-12, 2012, May 2012.

[Gont2013]

Gont, F., "Beta release of the SI6 Network's IPv6 Toolkit (help wanted!)", Message posted to the IPv6 Hackers mailing-list Message-ID: <51184548.3030105@si6networks.com>, 2013, <<https://lists.si6networks.com/pipermail/ipv6hackers/2013-February/000947.html>>.

[I-D.cooper-6man-ipv6-address-generation-privacy]

Cooper, A., Gont, F., and D. Thaler, "Privacy Considerations for IPv6 Address Generation Mechanisms", [draft-cooper-6man-ipv6-address-generation-privacy-00](#) (work in progress), July 2013.

[I-D.eddy-rfc793bis-04]

Eddy, W., "Transmission Control Protocol Specification", [draft-eddy-rfc793bis-04](#) (work in progress), August 2014.

[I-D.gont-6man-predictable-fragment-id]

Gont, F., "Security Implications of Predictable Fragment Identification Values", [draft-gont-6man-predictable-fragment-id-03](#) (work in progress), January 2013.

[I-D.gont-6man-stable-privacy-addresses]

Gont, F., "A method for Generating Stable Privacy-Enhanced Addresses with IPv6 Stateless Address Autoconfiguration (SLAAC)", [draft-gont-6man-stable-privacy-addresses-01](#) (work in progress), March 2012.

[I-D.gont-ntp-port-randomization]

Gont, F. and G. Gont, "Port Randomization in the Network Time Protocol Version 4", [draft-gont-ntp-port-randomization-04](#) (work in progress), August 2019.

[I-D.gont-numeric-ids-sec-considerations]

Gont, F. and I. Arce, "Security Considerations for Transient Numeric Identifiers Employed in Network Protocols", [draft-gont-numeric-ids-sec-considerations-06](#) (work in progress), December 2020.

[I-D.gont-opsec-ipv6-host-scanning]

Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", [draft-gont-opsec-ipv6-host-scanning-02](#) (work in progress), October 2012.

[I-D.gont-predictable-numeric-ids]

Gont, F. and I. Arce, "Security and Privacy Implications of Numeric Identifiers Employed in Network Protocols", [draft-gont-predictable-numeric-ids-03](#) (work in progress), March 2019.

[I-D.ietf-6man-default-iids]

Gont, F., Cooper, A., Thaler, D., and S. LIU, "Recommendation on Stable IPv6 Interface Identifiers", [draft-ietf-6man-default-iids-16](#) (work in progress), September 2016.

[I-D.ietf-6man-ipv6-address-generation-privacy]

Cooper, A., Gont, F., and D. Thaler, "Privacy Considerations for IPv6 Address Generation Mechanisms", [draft-ietf-6man-ipv6-address-generation-privacy-08](#) (work in progress), September 2015.

[I-D.ietf-6man-predictable-fragment-id]

Gont, F., "Security Implications of Predictable Fragment Identification Values", [draft-ietf-6man-predictable-fragment-id-10](#) (work in progress), October 2015.

[I-D.ietf-6man-rfc2460bis]

Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [draft-ietf-6man-rfc2460bis-13](#) (work in progress), May 2017.

[I-D.ietf-6man-rfc4941bis]

Gont, F., Krishnan, S., Narten, T., and R. Draves, "Temporary Address Extensions for Stateless Address Autoconfiguration in IPv6", [draft-ietf-6man-rfc4941bis-12](#) (work in progress), November 2020.

[I-D.ietf-6man-stable-privacy-addresses]

Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", [draft-ietf-6man-stable-privacy-addresses-17](#) (work in progress), January 2014.

[I-D.ietf-ntp-port-randomization]

Gont, F., Gont, G., and M. Lichvar, "Port Randomization in the Network Time Protocol Version 4", [draft-ietf-ntp-port-randomization-06](#) (work in progress), September 2020.

[I-D.ietf-ntp-refid-updates]

Stenn, H. and S. Goldberg, "Network Time Protocol REFID Updates", [draft-ietf-ntp-refid-updates-05](#) (work in progress), March 2019.

[I-D.ietf-opsec-ipv6-host-scanning]

Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", [draft-ietf-opsec-ipv6-host-scanning-08](#) (work in progress), August 2015.

[I-D.irtf-pearg-numeric-ids-generation]

Gont, F. and I. Arce, "On the Generation of Transient Numeric Identifiers", [draft-irtf-pearg-numeric-ids-generation-05](#) (work in progress), January 2021.

[IANA-PROT]

IANA, "Protocol Registries",
<<https://www.iana.org/protocols>>.

[IPID-DEV]

Klein, A. and B. Pinkas, "From IP ID to Device ID and KASLR Bypass (Extended Version)", June 2019,
<<https://arxiv.org/pdf/1906.10478.pdf>>.

[IPv6-Toolkit]

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

[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>.

[Morbitzer2013]

Morbitzer, M., "[PATCH] TCP Idle Scan in IPv6", Message posted to the nmap-dev mailing-list, 2013,
<<http://seclists.org/nmap-dev/2013/q2/394>>.

[Morris1985]

Morris, R., "A Weakness in the 4.2BSD UNIX TCP/IP Software", CSTR 117, AT&T Bell Laboratories, Murray Hill, NJ, 1985,
<<https://pdos.csail.mit.edu/~rtm/papers/117.pdf>>.

[NIST-NTP]

Sherman, J. and J. Levine, "Usage Analysis of the NIST Internet Time Service", Journal of Research of the National Institute of Standards and Technology Volume 121, March 2016, <<https://tf.nist.gov/general/pdf/2818.pdf>>.

[NTP-PORTR]

Gont, F., "[Ntp] New rev of the NTP port randomization I-D (Fwd: New Version Notification for [draft-gont-ntp-port-randomization-01.txt](#))", 2019, <<https://mailarchive.ietf.org/arch/browse/ntp/?glt=1&index=n09Sb61WkH031SRtamkELXwEQN4>>.

[OpenBSD-IPv4-ID]

OpenBSD, "Randomization of the IPv4 Identification field", December 1998, <https://cvsweb.openbsd.org/src/sys/netinet/ip_id.c?rev=1.1>.

[OpenBSD-IPv6-ID]

OpenBSD, "Randomization of the IPv6 Identification field", October 2003, <https://cvsweb.openbsd.org/src/sys/netinet6/ip6_id.c?rev=1.1>.

[OpenBSD-PR]

OpenBSD, "Implementation of port randomization", July 1996, <https://cvsweb.openbsd.org/src/sys/netinet/in_pcb.c?rev=1.6>.

[OpenBSD-TCP-ISN-H]

OpenBSD, "Implementation of [RFC1948](#) for TCP ISN randomization", December 2000, <https://cvsweb.openbsd.org/src/sys/netinet/tcp_subr.c?rev=1.97>.

[OpenBSD-TCP-ISN-I]

OpenBSD, "Implementation of TCP ISN randomization based on random increments", July 1996, <https://cvsweb.openbsd.org/src/sys/netinet/tcp_subr.c?rev=1.6>.

[OpenBSD-TCP-ISN-R]

OpenBSD, "Implementation of TCP ISN randomization based on simple randomization", December 2000, <https://cvsweb.openbsd.org/src/sys/netinet/tcp_subr.c?rev=1.37>.

[RedHat2011]

RedHat, "RedHat Security Advisory RHSA-2011:1465-1: Important: kernel security and bug fix update", 2011, <<https://rhn.redhat.com/errata/RHSA-2011-1465.html>>.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.

[RFC1948] Bellovin, S., "Defending Against Sequence Number Attacks", [RFC 1948](#), DOI 10.17487/RFC1948, May 1996, <<https://www.rfc-editor.org/info/rfc1948>>.

[RFC5157] Chown, T., "IPv6 Implications for Network Scanning", [RFC 5157](#), DOI 10.17487/RFC5157, March 2008, <<https://www.rfc-editor.org/info/rfc5157>>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", [RFC 5905](#), DOI 10.17487/RFC5905, June 2010, <<https://www.rfc-editor.org/info/rfc5905>>.

[RFC6274] Gont, F., "Security Assessment of the Internet Protocol Version 4", [RFC 6274](#), DOI 10.17487/RFC6274, July 2011, <<https://www.rfc-editor.org/info/rfc6274>>.

[RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", [BCP 188](#), [RFC 7258](#), DOI 10.17487/RFC7258, May 2014, <<https://www.rfc-editor.org/info/rfc7258>>.

[RFC7707] Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", [RFC 7707](#), DOI 10.17487/RFC7707, March 2016, <<https://www.rfc-editor.org/info/rfc7707>>.

[RFC7721] Cooper, A., Gont, F., and D. Thaler, "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", [RFC 7721](#), DOI 10.17487/RFC7721, March 2016, <<https://www.rfc-editor.org/info/rfc7721>>.

[RFC7739] Gont, F., "Security Implications of Predictable Fragment Identification Values", [RFC 7739](#), DOI 10.17487/RFC7739, February 2016, <<https://www.rfc-editor.org/info/rfc7739>>.

[RFC8064] Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", [RFC 8064](#), DOI 10.17487/RFC8064, February 2017, <<https://www.rfc-editor.org/info/rfc8064>>.

[Sanfilippo1998a]

Sanfilippo, S., "about the ip header id", Post to Bugtraq mailing-list, Mon Dec 14 1998,
<<http://seclists.org/bugtraq/1998/Dec/48>>.

[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>>.

[Schuba1993]

Schuba, C., "ADDRESSING WEAKNESSES IN THE DOMAIN NAME SYSTEM PROTOCOL", 1993,
<<http://ftp.cerias.purdue.edu/pub/papers/christoph-schuba/schuba-DNS-msthesis.pdf>>.

[Shimomura1995]

Shimomura, T., "Technical details of the attack described by Markoff in NYT", Message posted in USENET's comp.security.misc newsgroup Message-ID:
<3g5gk1\$5j1@ariel.sdsc.edu>, 1995,
<<https://www.gont.com.ar/docs/post-shimomura-usenet.txt>>.

[Silbersack2005]

Silbersack, M., "Improving TCP/IP security through randomization without sacrificing interoperability", EuroBSDCon 2005 Conference, 2005,
<<http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.91.4542&rep=rep1&type=pdf>>.

[SUSE2011]

SUSE, "SUSE Security Announcement: Linux kernel security update (SUSE-SA:2011:046)", 2011,
<<http://lists.opensuse.org/opensuse-security-announce/2011-12/msg00011.html>>.

[Ubuntu2011]

Ubuntu, "Ubuntu: USN-1253-1: Linux kernel vulnerabilities", 2011,
<<http://www.ubuntu.com/usn/usn-1253-1/>>.

[USCERT2001]

US-CERT, "US-CERT Vulnerability Note VU#498440: Multiple TCP/IP implementations may use statistically predictable initial sequence numbers", 2001,
<<https://www.kb.cert.org/vuls/id/498440>>.

[VU-800113]

CERT/CC, "Multiple DNS implementations vulnerable to cache poisoning (Vulnerability Note VU#800113)", July 2008,
<<https://www.kb.cert.org/vuls/id/800113>>.

[Wright1994]

Wright, G. and W. Stevens, "TCP/IP Illustrated, Volume 2: The Implementation", Addison-Wesley, 1994.

[Zalewski2001]

Zalewski, M., "Strange Attractors and TCP/IP Sequence Number Analysis", 2001,
<<http://lcamtuf.coredump.cx/oldtcp/tcpseq.html>>.

[Zalewski2002]

Zalewski, M., "Strange Attractors and TCP/IP Sequence Number Analysis - One Year Later", 2001,
<<http://lcamtuf.coredump.cx/newtcp/>>.

[Zalewski2003]

Zalewski, M., "A new TCP/IP blind data injection technique?", 2003,
<<http://lcamtuf.coredump.cx/ipfrag.txt>>.

Authors' Addresses

Fernando Gont
SI6 Networks
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

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

Ivan Arce
Quarkslab

Email: iarce@quarkslab.com
URI: <https://www.quarkslab.com>

