

Internet Research Task Force (IRTF)
Request for Comments: 9414
Category: Informational
ISSN: 2070-1721

F. Gont
SI6 Networks
I. Arce
Quarkslab
July 2023

Unfortunate History of Transient Numeric Identifiers

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 Enhancements and Assessments Research Group (PEARG) in the IRTF.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Research Task Force (IRTF). The IRTF publishes the results of Internet-related research and development activities. These results might not be suitable for deployment. This RFC represents the consensus of the Privacy Enhancements and Assessments Research Group of the Internet Research Task Force (IRTF). Documents approved for publication by the IRSG are not candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9414>.

Copyright Notice

Copyright (c) 2023 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. Terminology
3. Threat Model
4. Issues with the Specification of Transient Numeric Identifiers
 - 4.1. IPv4/IPv6 Identification
 - 4.2. TCP Initial Sequence Numbers (ISNs)
 - 4.3. IPv6 Interface Identifiers (IIDs)
 - 4.4. NTP Reference IDs (REFIDs)
 - 4.5. Transport-Protocol Ephemeral Port Numbers
 - 4.6. DNS ID
5. Conclusions

6.	IANA Considerations
7.	Security Considerations
8.	References
8.1.	Normative References
8.2.	Informative References
	Acknowledgements
	Authors' Addresses

1. Introduction

Networking protocols employ a variety of transient numeric identifiers for different protocol objects, such as IPv4 and IPv6 Identification values [RFC0791] [RFC8200], IPv6 Interface Identifiers (IIDs) [RFC4291], transport-protocol ephemeral port numbers [RFC6056], TCP Initial Sequence Numbers (ISNs) [RFC9293], NTP Reference IDs (REFIDs) [RFC5905], and DNS IDs [RFC1035]. These identifiers typically have specific requirements (e.g., uniqueness during a specified period of time) that must be satisfied such that they do not result in negative interoperability implications and an associated failure severity when such requirements are not met [RFC9415].

| NOTE: Some documents refer to the DNS ID as the DNS "Query ID"
| or "TxID".

For more than 30 years, a large number of implementations of IETF protocols 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, the poor selection of transient numeric identifiers in such protocols, usually as a result of insufficient or misleading specifications. While it is generally trivial to identify an algorithm that can satisfy the interoperability requirements of a given transient numeric identifier, empirical evidence exists that doing so without negatively affecting the security and/or privacy properties of the aforementioned protocols is prone to error.

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

- * predictable IPv4 or IPv6 Identification values (e.g., see [Sanfilippo1998a], [RFC6274], and [RFC7739]),
- * predictable IPv6 IIDs (e.g., see [RFC7217], [RFC7707], and [RFC7721]),
- * predictable transport-protocol ephemeral port numbers (e.g., see [RFC6056] and [Silbersack2005]),
- * predictable TCP Initial Sequence Numbers (ISNs) (e.g., see [Morris1985], [Bellovin1989], and [RFC6528]),
- * predictable initial timestamps in TCP timestamps options (e.g., see [TCPT-uptime] and [RFC7323]), and
- * predictable DNS IDs (see, e.g., [Schuba1993] and [Klein2007]).

Recent history indicates that, when new protocols are standardized or new protocol implementations are produced, the security and privacy properties of the associated transient numeric identifiers tend to be overlooked, and inappropriate algorithms to generate such identifiers are either suggested in the specifications or selected by implementers. As a result, advice in this area is warranted.

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:

- * vulnerabilities associated with the inappropriate generation of transient numeric identifiers have affected protocol implementations for an extremely long period of time,
- * such vulnerabilities, even when addressed for a given protocol version, were later reintroduced in new versions or new implementations of the same protocol, and
- * standardization efforts that discuss and provide advice in this area can have a positive effect on IETF specifications and their corresponding 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 and/or privacy properties of the corresponding protocols is nontrivial. Other related documents ([RFC9415] and [RFC9416]) provide guidance in this area, as motivated by the present document.

This document represents the consensus of the Privacy Enhancements and Assessments 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 (e.g., see [IANA-PROT]). We note that different transient numeric 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 do not consider on-path attacks. That is, we assume the attacker does not have physical or logical access to the system(s) being attacked and that the attacker can only observe traffic explicitly directed to the attacker. Similarly, an attacker cannot observe traffic transferred between the sender and the receiver(s) of a target protocol but may be able to interact with any of these entities, including by, e.g., sending any traffic to them to sample transient numeric identifiers employed by the target hosts when communicating with the attacker.

For example, when analyzing vulnerabilities associated with TCP Initial Sequence Numbers (ISNs), we consider the attacker is unable to capture network traffic corresponding to a TCP connection between two other hosts. However, we consider the attacker is able to communicate with any of these hosts (e.g., establish a TCP connection with any of them) to, e.g., sample the TCP ISNs employed by these

hosts when communicating with the attacker.

Similarly, when considering host-tracking attacks based on IPv6 Interface Identifiers, we consider an attacker may learn the IPv6 address employed by a victim host if, e.g., the address becomes exposed as a result of the victim host communicating with an attacker-operated server. Subsequently, an attacker may perform host-tracking by probing a set of target addresses composed by a set of target prefixes and the IPv6 Interface Identifier originally learned by the attacker. Alternatively, an attacker may perform host-tracking if, e.g., the victim host communicates with an attacker-operated server as it moves from one location to another, thereby exposing its configured addresses. We note that none of these scenarios require the attacker observe traffic not explicitly directed to the attacker.

4. Issues with the Specification of Transient Numeric Identifiers

While assessing IETF 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:

- * protocol specifications that under specify their transient numeric identifiers
- * protocol specifications that over specify their transient numeric identifiers
- * protocol implementations that simply fail to comply with the specified requirements

A number of IETF protocol specifications under specified their transient numeric identifiers, thus leading to implementations that were vulnerable to numerous off-path attacks. Examples of them are the specification of TCP local ports in [RFC0793] or the specification of the DNS ID in [RFC1035].

| NOTE: The TCP local port in an active OPEN request is commonly
| known as the "ephemeral port" of the corresponding TCP
| connection [RFC6056].

On the other hand, there are a number of IETF 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] suggests the use of a global counter for the generation of Identification values when the interoperability requirement of uniqueness per {IPv6 Source Address, IPv6 Destination Address} 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 still fail to implement transport-protocol ephemeral port randomization, as recommended in [RFC6056], or TCP Initial Sequence Number randomization, as recommended in [RFC9293].

The following subsections document the timelines for a number of sample transient numeric identifiers that illustrate how the problem discussed in this document has affected protocols from different layers over time. These sample transient numeric identifiers have different interoperability requirements and failure severities (see

Section 6 of [RFC9415]), and thus are considered to be representative of the problem being analyzed in this document.

4.1. 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 protocols, the work and research done for the IPv4 case did not influence IPv6 specifications or implementations.

The IPv4 Identification is specified in [RFC0791], which specifies the interoperability requirements for the Identification field, i.e., the sender must choose the Identification field to be unique for a given {Source Address, Destination Address, Protocol} for the time the datagram (or any fragment of it) could be alive in the Internet. It suggests that a sending protocol module may keep "a table of Identifiers, one entry for each destination it has communicated with in the last maximum packet lifetime for the [I]nternet", and it also 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 resulted in many IPv4 implementations generating predictable 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, but rather than being part of the base header (as in the IPv4 case), it is part of the Fragment Header (which may or may not be present in an IPv6 packet). Section 4.5 of [RFC1883] states 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 specification 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] is the first revision of the core IPv6 specification and maintains 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 the IPv4 Identification 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 values 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} set).

December 1998:

[Sanfilippo1998a] finds that predictable IPv4 Identification values (as 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 "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].

December 1998:

OpenBSD implements randomization of the IPv4 Identification field [OpenBSD-IPv4-ID].

November 1999:

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

September 2002:

[Fyodor2002] documents the implementation of the "idle scan" technique in the popular Network Mapper (nmap) tool.

November 2002:

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

October 2003:

OpenBSD implements randomization of the IPv6 Identification field [OpenBSD-IPv6-ID].

December 2003:

[Zalewski2003] explains a technique to perform TCP data injection attacks based on predictable IPv4 Identification values, which requires less effort than TCP injection attacks performed with bare TCP packets.

January 2005:

[Silbersack2005] discusses shortcomings in a number of techniques to mitigate predictable IPv4 Identification values.

October 2007:

[Klein2007] describes a weakness in the pseudorandom number generator (PRNG) in use for the generation of IP Identification values 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 the generation of IPv6 Identification values.

May 2012:

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

March 2013:

The 6man WG adopts [draft-gont-6man-predictable-fragment-id-03] but changes the track to "BCP" (while still formally updating [RFC2460]), posting the resulting document as [draft-ietf-6man-predictable-fragment-id-00].

June 2013:

A patch to incorporate support for IPv6-based idle scans 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 posts it as [draft-ietf-6man-predictable-fragment-id-02]. As a result, it no longer formally updates [RFC2460], and security and privacy requirements on the generation of IPv6 Identification values are eliminated.

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 [draft-ietf-6man-predictable-fragment-id-10]) 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 as an "Informational" RFC, it does not formally update [RFC2460] and does not introduce security and privacy requirements on the generation of IPv6 Identification values.

June 2016:

[draft-ietf-6man-rfc2460bis-05], a draft 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 and does not introduce security and privacy requirements on the generation of IPv6 Identification values.

July 2017:

[draft-ietf-6man-rfc2460bis-13] is finally published as [RFC8200], obsoleting [RFC2460] and pointing to [RFC7739] for sample algorithms for the generation of IPv6 Identification values. However, it does not introduce security and privacy requirements on the generation of IPv6 Identification values.

October 2019:

[IPID-DEV] notes that the IPv6 Identification generators of two popular operating systems are flawed.

4.2. 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 $2 \times \text{MSL}$ has 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 implementation will use for new connections, thus allowing a variety of attacks against TCP.

February 1985:

[Morris1985] is 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] discusses the security considerations for predictable ISNs (along with a range of other protocol-based vulnerabilities).

January 1995:

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

May 1996:

[RFC1948] is the first IETF effort, authored by Steven Bellovin, to address predictable TCP ISNs. However, [RFC1948] does not formally update [RFC0793]. Note: 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 [RFC9415]) [OpenBSD-TCP-ISN-I].

December 2000:

OpenBSD implements TCP ISN randomization using simple randomization (please see Section 7.1 of [RFC9415]) [OpenBSD-TCP-ISN-R].

March 2001:

[Zalewski2001] provides a detailed analysis of statistical weaknesses in some TCP ISN generators and includes a survey of the algorithms in use by popular TCP implementations. Vulnerability advisories [USCERT2001] were released regarding statistical weaknesses in some TCP ISN generators, affecting popular TCP implementations. Other vulnerability advisories on the same vulnerability, such as [CERT2001], were published later on.

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 and replaced 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".

February 2012:

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

August 2014:

The algorithm specified in [RFC6528] becomes the recommended ("SHOULD") algorithm for TCP ISN generation in [draft-eddy-rfc793bis-04], an early revision of the core TCP specification [RFC9293].

August 2022:

[RFC9293], a revision of the core TCP specification, is published, adopting the algorithm specified in [RFC6528] as the recommended ("SHOULD") algorithm for TCP ISN generation.

4.3. IPv6 Interface Identifiers (IIDs)

IPv6 Interface Identifiers can be generated as a result of different mechanisms, including Stateless Address Autoconfiguration (SLAAC) [RFC4862], DHCPv6 [RFC8415], and manual configuration. This section focuses on Interface Identifiers resulting from SLAAC.

The Interface Identifier of stable IPv6 addresses resulting from SLAAC originally 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), which is 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 [RFC2073]), changing the size of the IID to 64 bits, and specifies that IIDs must be constructed in IEEE 64-bit Extended Unique Identifier (EUI-64) format. How such identifiers are constructed is specified in the corresponding "IPv6 over <link>" specifications, such as "IPv6 over Ethernet".

January 2001:

[RFC3041] recognizes the problem of IPv6 network activity correlation and specifies IPv6 temporary addresses. Temporary addresses are to be used along with stable addresses.

August 2003:

[RFC3587] obsoletes [RFC2374], making the Top-Level Aggregator (TLA) / Next-Level Aggregator (NLA) structure historic, though the syntax and recommendations for the stable IIDs remain unchanged.

February 2006:

[RFC4291] is published as the latest "IP Version 6 Addressing Architecture", requiring the IIDs of "all unicast addresses, except those that start with the binary value 000" to employ the Modified EUI-64 format. The details of constructing such interface identifiers are defined in the corresponding "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 original scheme for generating stable addresses allows for IPv6 address scanning and for active host tracking (even when IPv6 temporary addresses are employed). It also specifies an alternative algorithm meant to replace IIDs based on Modified EUI-64 format identifiers.

November 2012:

The 6man WG adopts [draft-gont-6man-stable-privacy-addresses-01] as a working group item (as [draft-ietf-6man-stable-privacy-addresses-00]). However, the document no longer formally updates [RFC4291]; 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:

[draft-cooper-6man-ipv6-address-generation-privacy-00] elaborates on the security and privacy properties of all known algorithms for generating IPv6 IIDs.

January 2014:

The 6man WG posts [draft-ietf-6man-default-iids-00] ("Recommendation on Stable IPv6 Interface Identifiers"), recommending [draft-ietf-6man-stable-privacy-addresses-17] for the generation of stable addresses.

April 2014:

[RFC7217] (formerly [draft-ietf-6man-stable-privacy-addresses-17]) 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 [draft-gont-opsec-ipv6-host-scanning-02] and later [draft-ietf-opsec-ipv6-host-scanning-08]), about "Network Reconnaissance in IPv6 Networks", is published.

March 2016:

[RFC7721] (formerly [draft-cooper-6man-ipv6-address-generation-privacy-00] and later [draft-ietf-6man-ipv6-address-generation-privacy-08]), about "Security and Privacy Considerations for IPv6 Address Generation Mechanisms", is published.

May 2016:

[draft-gont-6man-non-stable-iids-00] is posted, 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 posted, 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:

[draft-ietf-6man-default-iids-16], produced by the 6man WG, is published as [RFC8064] ("Recommendation on Stable IPv6 Interface Identifiers"), 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 posted (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, posted in March 2016).

July 2018:

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

December 2020:

[draft-ietf-6man-rfc4941bis-12] is approved by the IESG for publication as an RFC.

February 2021:

[draft-ietf-6man-rfc4941bis-12] is finally published as [RFC8981].

4.4. NTP Reference IDs (REFIDs)

The NTP [RFC5905] Reference ID is a 32-bit code identifying the particular server or reference clock. Above stratum 1 (secondary servers and clients), this value can be employed to avoid degree-one timing loops, that is, scenarios where two NTP peers are (mutually) the time source of each other. If using the IPv4 address family, the identifier is the four-octet IPv4 address. If using the IPv6 address family, it is the first four octets of the MD5 hash of the IPv6 address.

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 the first four octets of the MD5 hash of the IPv6 address of the time source.

July 2016:

[draft-stenn-ntp-not-you-refid-00] is posted, 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 common REFID. It is

subsequently adopted by the NTP WG as [draft-ietf-ntp-refid-updates-00].

April 2019:

[Gont-NTP] notes that the proposed fix specified in [draft-ietf-ntp-refid-updates-00] is, at the very least, sub-optimal. As a result of a lack of WG support, the [draft-ietf-ntp-refid-updates-00] effort is eventually abandoned.

4.5. Transport-Protocol Ephemeral Port Numbers

Most (if not all) transport protocols employ "port numbers" to demultiplex packets to the corresponding transport-protocol instances. "Ephemeral ports" refer to the local ports employed in active OPEN requests, that is, typically the local port numbers employed on the side initiating the communication.

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 are selected and the port range from which they are selected are left unspecified.

July 1996:

OpenBSD implements ephemeral port randomization [OpenBSD-PR].

July 2008:

The CERT Coordination Center publishes details of what became known as the "Kaminsky Attack" [VU-800113] [Kaminsky2008] on the DNS. The attack exploits the lack of ephemeral port randomization and DNS ID randomization 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 of available ports (53 or 1024 and above) that is as large as possible and practicable. 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 resolvers themselves, rather than in the underlying transport protocols.

January 2011:

[RFC6056] notes that many TCP and UDP implementations result in predictable ephemeral 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:

[draft-gont-ntp-port-randomization-00] notes that some NTP implementations employ the NTP service port (123) as the local port for nonsymmetric modes and aims to update the NTP specification to recommend port randomization in such cases, which is in line with [RFC6056]. The proposal experiences some pushback in the relevant working group (NTP WG) [NTP-PORTR] but is finally adopted as a working group item as [draft-ietf-ntp-port-randomization-00].

August 2021:

[draft-ietf-ntp-port-randomization-08] is finally published as [RFC9109].

4.6. DNS ID

The DNS ID [RFC1035] can be employed to match DNS replies to outstanding DNS queries.

| NOTE: Some documents refer to the DNS ID as the DNS "Query ID"
| or "TxID".

November 1987:

[RFC1035] specifies that the DNS ID is a 16-bit identifier assigned by the program that generates any kind of query and that this identifier is copied in the corresponding reply and can be used by the requester to match up replies to outstanding queries. It does not specify the interoperability requirements for this numeric identifier, nor does it suggest an algorithm for generating it.

August 1993:

[Schubal1993] describes DNS cache poisoning attacks that require the attacker to guess the DNS ID.

June 1995:

[Vixiel1995] suggests that both the UDP source port and the DNS ID of query packets should be randomized, although that might not provide enough entropy to prevent an attacker from guessing these values.

April 1997:

[Arcel1997] finds that implementations employ predictable UDP source ports and predictable DNS IDs and argues that both should be randomized.

November 2002:

[Sacramento2002] finds that, by spoofing multiple requests for the same domain name from different IP addresses, an attacker may guess the DNS ID employed for a victim with a high probability of success, thus allowing for DNS cache poisoning attacks.

March 2007:

[Klein2007c] finds that the Microsoft Windows DNS server generates predictable DNS ID values.

July 2007:

[Klein2007b] finds that a popular DNS server software (BIND 9) that randomizes the DNS ID is still subject to DNS cache poisoning attacks by forging a large number of queries and leveraging the birthday paradox.

October 2007:

[Klein2007] finds that OpenBSD's DNS software (based on the BIND DNS server of the Internet Systems Consortium (ISC)) generates

predictable DNS ID values.

January 2009:

[RFC5452] is published, requiring resolvers to randomize the DNS ID of queries and to verify that the DNS ID of a reply matches that of the DNS query as part of the DNS reply validation process.

May 2010:

[Economou2010] finds that the Windows SMTP Service implements its own DNS resolver that results in predictable DNS ID values. Additionally, it fails to validate that the DNS ID of a reply matches that of the DNS query that supposedly elicited it.

5. Conclusions

For more than 30 years, a large number of implementations of IETF protocols 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, the poor selection of transient numeric identifiers in such protocols, usually as a result of insufficient or misleading specifications.

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 and/or privacy properties of the aforementioned protocols is nontrivial. It is thus evident that advice in this area is warranted.

[RFC9416] aims at requiring future IETF protocol specifications to contain analysis of the security and privacy properties of any transient numeric identifiers specified by the protocol and to recommend an algorithm for the generation of such transient numeric identifiers. [RFC9415] specifies a number of sample algorithms for generating transient numeric identifiers with specific interoperability requirements and failure severities.

6. IANA Considerations

This document has no IANA actions.

7. 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.

[RFC9415] analyzes and categorizes transient numeric identifiers based on their interoperability requirements and their associated failure severities and recommends possible algorithms that can be employed to comply with those requirements without negatively affecting the security and privacy properties of the corresponding protocols.

[RFC9416] formally requires IETF 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 or privacy properties of these transient numeric identifiers are mitigated.

8. References

8.1. 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", 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>>.
- [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>>.
- [RFC7323] Borman, D., Braden, B., Jacobson, V., and R. Scheffenegger, Ed., "TCP Extensions for High Performance", RFC 7323, DOI 10.17487/RFC7323, September 2014, <<https://www.rfc-editor.org/info/rfc7323>>.
- [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>>.
- [RFC9293] Eddy, W., Ed., "Transmission Control Protocol (TCP)", STD 7, RFC 9293, DOI 10.17487/RFC9293, August 2022, <<https://www.rfc-editor.org/info/rfc9293>>.

8.2. Informative References

- [Arce1997] Arce, I. and E. Kargieman, "BIND Vulnerabilities and Solutions", April 1997, <http://www.openbsd.org/advisories/sni_12_resolverid.txt>.
- [Bellovin1989] Bellovin, S., "Security Problems in the TCP/IP Protocol Suite", Computer Communications Review, vol. 19, no. 2, pp. 32-48, April 1989, <<https://www.cs.columbia.edu/~smb/papers/ipext.pdf>>.
- [Bellovin2002] Bellovin, S., "A Technique for Counting NATted Hosts", IMW'02, Marseille, France, DOI 10.1145/637201.637243, November 2002, <<https://www.cs.columbia.edu/~smb/papers/fnat.pdf>>.
- [CERT2001] CERT/CC, "CERT Advisory CA-2001-09: Statistical Weaknesses in TCP/IP Initial Sequence Numbers", May 2001, <https://resources.sei.cmu.edu/asset_files/WhitePaper/2001_019_001_496192.pdf>.
- [draft-cooper-6man-ipv6-address-generation-privacy-00] Cooper, A., Gont, F., and D. Thaler, "Privacy Considerations for IPv6 Address Generation Mechanisms", Work in Progress, Internet-Draft, draft-cooper-6man-ipv6-address-generation-privacy-00, 15 July 2013,

<<https://www.ietf.org/archive/id/draft-cooper-6man-ipv6-address-generation-privacy-00.txt>>.

[draft-eddy-rfc793bis-04]

Eddy, W., Ed., "Transmission Control Protocol Specification", Work in Progress, Internet-Draft, draft-eddy-rfc793bis-04, 25 August 2014, <<https://www.ietf.org/archive/id/draft-eddy-rfc793bis-04.txt>>.

[draft-fgont-6man-rfc4941bis-00]

Gont, F., Krishnan, S., Narten, T., and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", Work in Progress, Internet-Draft, draft-fgont-6man-rfc4941bis-00, 25 March 2018, <<https://www.ietf.org/archive/id/draft-fgont-6man-rfc4941bis-00.txt>>.

[draft-gont-6man-address-usage-recommendations-00]

Gont, F. and W. LIU, "IPv6 Address Usage Recommendations", Work in Progress, Internet-Draft, draft-gont-6man-address-usage-recommendations-00, 27 May 2016, <<https://www.ietf.org/archive/id/draft-gont-6man-address-usage-recommendations-00.txt>>.

[draft-gont-6man-non-stable-iids-00]

Gont, F. and W. Liu, "Recommendation on Non-Stable IPv6 Interface Identifiers", Work in Progress, Internet-Draft, draft-gont-6man-non-stable-iids-00, 23 May 2016, <<https://www.ietf.org/archive/id/draft-gont-6man-non-stable-iids-00.txt>>.

[draft-gont-6man-predictable-fragment-id-00]

Gont, F., "Security Implications of Predictable Fragment Identification Values", Work in Progress, Internet-Draft, draft-gont-6man-predictable-fragment-id-00, 15 December 2011, <<https://www.ietf.org/archive/id/draft-gont-6man-predictable-fragment-id-00.txt>>.

[draft-gont-6man-predictable-fragment-id-03]

Gont, F., "Security Implications of Predictable Fragment Identification Values", Work in Progress, Internet-Draft, draft-gont-6man-predictable-fragment-id-03, 9 January 2013, <<https://www.ietf.org/archive/id/draft-gont-6man-predictable-fragment-id-03.txt>>.

[draft-gont-6man-stable-privacy-addresses-00]

Gont, F., "A method for Generating Stable Privacy-Enhanced Addresses with IPv6 Stateless Address Autoconfiguration (SLAAC)", Work in Progress, Internet-Draft, draft-gont-6man-stable-privacy-addresses-00, 15 December 2011, <<https://www.ietf.org/archive/id/draft-gont-6man-stable-privacy-addresses-00.txt>>.

[draft-gont-6man-stable-privacy-addresses-01]

Gont, F., "A method for Generating Stable Privacy-Enhanced Addresses with IPv6 Stateless Address Autoconfiguration (SLAAC)", Work in Progress, Internet-Draft, draft-gont-6man-stable-privacy-addresses-01, 31 March 2012, <<https://www.ietf.org/archive/id/draft-gont-6man-stable-privacy-addresses-01.txt>>.

[draft-gont-ntp-port-randomization-00]

Gont, F. and G. Gont, "Port Randomization in the Network Time Protocol Version 4", Work in Progress, Internet-Draft, draft-gont-ntp-port-randomization-00, 16 April

2019, <<https://www.ietf.org/archive/id/draft-gont-ntp-port-randomization-00.txt>>.

[draft-gont-opsec-ipv6-host-scanning-02]

Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", Work in Progress, Internet-Draft, draft-gont-opsec-ipv6-host-scanning-02, 23 October 2012, <<https://www.ietf.org/archive/id/draft-gont-opsec-ipv6-host-scanning-02.txt>>.

[draft-gont-predictable-numeric-ids-03]

Gont, F. and I. Arce, "Security and Privacy Implications of Numeric Identifiers Employed in Network Protocols", Work in Progress, Internet-Draft, draft-gont-predictable-numeric-ids-03, 11 March 2019, <<https://datatracker.ietf.org/doc/html/draft-gont-predictable-numeric-ids-03>>.

[draft-ietf-6man-default-iids-00]

Gont, F., Cooper, A., Thaler, D., and W. Liu, "Recommendation on Stable IPv6 Interface Identifiers", Work in Progress, Internet-Draft, draft-ietf-6man-default-iids-00, 24 January 2014, <<https://www.ietf.org/archive/id/draft-ietf-6man-default-iids-00.txt>>.

[draft-ietf-6man-default-iids-16]

Gont, F., Cooper, A., Thaler, D., and W. LIU, "Recommendation on Stable IPv6 Interface Identifiers", Work in Progress, Internet-Draft, draft-ietf-6man-default-iids-16, 28 September 2016, <<https://www.ietf.org/archive/id/draft-ietf-6man-default-iids-16.txt>>.

[draft-ietf-6man-ipv6-address-generation-privacy-08]

Cooper, A., Gont, F., and D. Thaler, "Privacy Considerations for IPv6 Address Generation Mechanisms", Work in Progress, Internet-Draft, draft-ietf-6man-ipv6-address-generation-privacy-08, 23 September 2015, <<https://www.ietf.org/archive/id/draft-ietf-6man-ipv6-address-generation-privacy-08.txt>>.

[draft-ietf-6man-predictable-fragment-id-00]

Gont, F., "Security Implications of Predictable Fragment Identification Values", Work in Progress, Internet-Draft, draft-ietf-6man-predictable-fragment-id-00, 22 March 2013, <<https://www.ietf.org/archive/id/draft-ietf-6man-predictable-fragment-id-00.txt>>.

[draft-ietf-6man-predictable-fragment-id-01]

Gont, F., "Security Implications of Predictable Fragment Identification Values", Work in Progress, Internet-Draft, draft-ietf-6man-predictable-fragment-id-01, 29 April 2014, <<https://www.ietf.org/archive/id/draft-ietf-6man-predictable-fragment-id-01.txt>>.

[draft-ietf-6man-predictable-fragment-id-02]

Gont, F., "Security Implications of Predictable Fragment Identification Values", Work in Progress, Internet-Draft, draft-ietf-6man-predictable-fragment-id-02, 19 December 2014, <<https://datatracker.ietf.org/doc/html/draft-ietf-6man-predictable-fragment-id-02>>.

[draft-ietf-6man-predictable-fragment-id-08]

Gont, F., "Security Implications of Predictable Fragment Identification Values", Work in Progress, Internet-Draft,

draft-ietf-6man-predictable-fragment-id-08, 9 June 2015,
<<https://www.ietf.org/archive/id/draft-ietf-6man-predictable-fragment-id-08.txt>>.

[draft-ietf-6man-predictable-fragment-id-10]

Gont, F., "Security Implications of Predictable Fragment Identification Values", Work in Progress, Internet-Draft, draft-ietf-6man-predictable-fragment-id-10, 9 October 2015, <<https://www.ietf.org/archive/id/draft-ietf-6man-predictable-fragment-id-10.txt>>.

[draft-ietf-6man-rfc2460bis-05]

Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", Work in Progress, Internet-Draft, draft-ietf-6man-rfc2460bis-05, 28 June 2016, <<https://www.ietf.org/archive/id/draft-ietf-6man-rfc2460bis-05.txt>>.

[draft-ietf-6man-rfc2460bis-13]

Deering, S. and R. Hinden, "draft-ietf-6man-rfc2460bis-13", Work in Progress, Internet-Draft, draft-ietf-6man-rfc2460bis-13, 19 May 2017, <<https://www.ietf.org/archive/id/draft-ietf-6man-rfc2460bis-13.txt>>.

[draft-ietf-6man-rfc4941bis-00]

Gont, F., Krishnan, S., Narten, T., and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", Work in Progress, Internet-Draft, draft-ietf-6man-rfc4941bis-00, 2 July 2018, <<https://www.ietf.org/archive/id/draft-ietf-6man-rfc4941bis-00.txt>>.

[draft-ietf-6man-rfc4941bis-12]

Gont, F., Krishnan, S., Narten, T., and R. Draves, "Temporary Address Extensions for Stateless Address Autoconfiguration in IPv6", Work in Progress, Internet-Draft, draft-ietf-6man-rfc4941bis-12, 2 November 2020, <<https://www.ietf.org/archive/id/draft-ietf-6man-rfc4941bis-12.txt>>.

[draft-ietf-6man-stable-privacy-addresses-00]

Gont, F., "A method for Generating Stable Privacy-Enhanced Addresses with IPv6 Stateless Address Autoconfiguration (SLAAC)", Work in Progress, Internet-Draft, draft-ietf-6man-stable-privacy-addresses-00, 18 May 2012, <<https://www.ietf.org/archive/id/draft-ietf-6man-stable-privacy-addresses-00.txt>>.

[draft-ietf-6man-stable-privacy-addresses-17]

Gont, F., "A Method for Generating Semantically Opaque Interface Identifiers with IPv6 Stateless Address Autoconfiguration (SLAAC)", Work in Progress, Internet-Draft, draft-ietf-6man-stable-privacy-addresses-17, 27 January 2014, <<https://www.ietf.org/archive/id/draft-ietf-6man-stable-privacy-addresses-17.txt>>.

[draft-ietf-ntp-port-randomization-00]

Gont, F., Gont, G., and M. Lichvar, "Port Randomization in the Network Time Protocol Version 4", Work in Progress, Internet-Draft, draft-ietf-ntp-port-randomization-00, 22 October 2019, <<https://www.ietf.org/archive/id/draft-ietf-ntp-port-randomization-00.txt>>.

[draft-ietf-ntp-port-randomization-08]

Gont, F., Gont, G., and M. Lichvar, "Port Randomization in

the Network Time Protocol Version 4", Work in Progress, Internet-Draft, draft-ietf-ntp-port-randomization-00, 10 June 2021, <<https://www.ietf.org/archive/id/draft-ietf-ntp-port-randomization-08.txt>>.

[draft-ietf-ntp-refid-updates-00]

Stenn, H. and S. Goldberg, "Network Time Protocol REFID Updates", Work in Progress, Internet-Draft, draft-ietf-ntp-refid-updates-00, 13 November 2016, <<https://www.ietf.org/archive/id/draft-ietf-ntp-refid-updates-00.txt>>.

[draft-ietf-opsec-ipv6-host-scanning-08]

Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", Work in Progress, Internet-Draft, draft-ietf-opsec-ipv6-host-scanning-08, 28 August 2015, <<https://www.ietf.org/archive/id/draft-ietf-opsec-ipv6-host-scanning-08.txt>>.

[draft-stenn-ntp-not-you-refid-00]

Goldberg, S. and H. Stenn, "Network Time Protocol Not You REFID", Work in Progress, Internet-Draft, draft-stenn-ntp-not-you-refid-00, 8 July 2016, <<https://www.ietf.org/archive/id/draft-stenn-ntp-not-you-refid-00.txt>>.

[Economou2010]

Economou, N., "Windows SMTP Service DNS query Id vulnerabilities", Advisory ID Internal CORE-2010-0427, May 2010, <<https://www.coresecurity.com/core-labs/advisories/core-2010-0424-windows-smtp-dns-query-id-bugs>>.

[Fyodor2002]

Fyodor, "Idle scanning and related IP ID games", September 2002, <https://nmap.org/presentations/CanSecWest03/CD_Content/idlescan_paper/idlescan.html>.

[Gont-NTP] Gont, F., "[Ntp] Comments on draft-ietf-ntp-refid-updates-05", message to the IETF NTP mailing list, 16 April 2019, <<https://mailarchive.ietf.org/arch/msg/ntp/NkfTHxUUOdp14Agh3h1IPqfcRRg>>.

[Gont2011] Gont, F., "Hacking IPv6 Networks (training course)", Hack In Paris 2011 Conference, Paris, France, June 2011, <<https://www.si6networks.com/files/presentations/hip2011/fgont-hip2011-hacking-ipv6-networks.pdf>>.

[Gont2012] Gont, F., "Recent Advances in IPv6 Security", BSDCan 2012 Conference, Ottawa, Canada, May 2012, <<https://www.si6networks.com/files/presentations/bsdcan2012/fgont-bsdcan2012-recent-advances-in-ipv6-security.pdf>>.

[Gont2013] Gont, F., "Beta release of the SI6 Network's IPv6 Toolkit (help wanted!)", message to the IPv6 Hackers mailing list, 11 February 2013, <<https://lists.si6networks.com/pipermail/ipv6hackers/2013-February/000947.html>>.

[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)",

DOI 10.48550/arXiv.1906.10478, October 2019,
<<https://arxiv.org/pdf/1906.10478.pdf>>.

[IPv6-Toolkit]

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

[Kaminsky2008]

Kaminsky, D., "Black Ops 2008: It's The End Of The Cache As We Know It", August 2008,
<<https://www.blackhat.com/presentations/bh-jp-08/bh-jp-08-Kaminsky/BlackHat-Japan-08-Kaminsky-DNS08-BlackOps.pdf>>.

[Klein2007]

Klein, A., "OpenBSD DNS Cache Poisoning and Multiple O/S Predictable IP ID Vulnerability", October 2007,
<https://dl.packetstormsecurity.net/papers/attack/OpenBSD_DNS_Cache_Poisoning_and_Multiple_OS_Predictable_IP_ID_Vulnerability.pdf>.

[Klein2007b]

Klein, A., "BIND 9 DNS Cache Poisoning", March 2007,
<<https://citeseerx.ist.psu.edu/doc/10.1.1.86.4474>>.

[Klein2007c]

Klein, A., "Windows DNS Server Cache Poisoning", March 2007, <https://dl.packetstormsecurity.net/papers/attack/Windows_DNS_Cache_Poisoning.pdf>.

[Morbitzer2013]

Morbitzer, M., "[PATCH] TCP Idle Scan in IPv6", message to the nmap-dev mailing list, 3 June 2013,
<<https://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, February 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)", message to the IETF NTP mailing list, 21 May 2019,
<<https://mailarchive.ietf.org/arch/msg/ntp/xSCu5Vhb3zoWcqEjUMmzP8pOdW4/>>.

[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", June 2007, <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]

Red Hat, "RHSA-2011:1465-1 - Security Advisory", November 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>>.
- [RFC8981] Gont, F., Krishnan, S., Narten, T., and R. Draves, "Temporary Address Extensions for Stateless Address Autoconfiguration in IPv6", RFC 8981, DOI 10.17487/RFC8981, February 2021, <<https://www.rfc-editor.org/info/rfc8981>>.
- [RFC9109] Gont, F., Gont, G., and M. Lichvar, "Network Time Protocol Version 4: Port Randomization", RFC 9109, DOI 10.17487/RFC9109, August 2021, <<https://www.rfc-editor.org/info/rfc9109>>.
- [RFC9415] Gont, F. and I. Arce, "On the Generation of Transient Numeric Identifiers", RFC 9415, DOI 10.17487/RFC9415, July 2023, <<https://www.rfc-editor.org/info/rfc9415>>.
- [RFC9416] Gont, F. and I. Arce, "Security Considerations for Transient Numeric Identifiers Employed in Network Protocols", BCP 72, RFC 9416, DOI 10.17487/RFC9416, July 2023, <<https://www.rfc-editor.org/info/rfc9416>>.
- [Sacramento2002] Sacramento, V., "CAIS-ALERT: Vulnerability in the sending requests control of BIND", message to the Bugtraq mailing list, 25 November 2002, <<https://seclists.org/bugtraq/2002/Nov/331>>.
- [Sanfilippo1998a] Sanfilippo, S., "about the ip header id", message to the Bugtraq mailing list, 14 December 1998, <<http://seclists.org/bugtraq/1998/Dec/48>>.
- [Sanfilippo1998b] Sanfilippo, S., "new tcp scan method", message to the Bugtraq mailing list, 18 December 1998, <<https://seclists.org/bugtraq/1998/Dec/79>>.
- [Sanfilippo1999] Sanfilippo, S., "more ip id", message to the Bugtraq mailing list, November 1999, <<https://github.com/antirez/hping/raw/master/docs/MORE-FUN-WITH-IPID>>.
- [Schuba1993] Schuba, C., "Addressing Weakness in the Domain Name System Protocol", August 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 to the USENET comp.security.misc newsgroup, 25 January 1995, <<https://www.gont.com.ar/files/post-shimomura-usenet.txt>>.
- [Silbersack2005] Silbersack, M., "Improving TCP/IP security through randomization without sacrificing interoperability", EuroBSDCon 2005 Conference, January 2005, <<https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.91.4542&rep=rep1&type=pdf>>.
- [SUSE2011] Meissner, M., "[security-announce] SUSE Security

Announcement: Linux kernel security update (SUSE-SA:2011:046)", message to the openSUSE Security Announce mailing list, 13 December 2011, <<https://lists.opensuse.org/opensuse-security-announce/2011-12/msg00011.html>>.

[TCPT-uptime]

McDanel, B., "TCP Timestamping - Obtaining System Uptime Remotely", message to the Bugtraq mailing list, March 2001, <<https://seclists.org/bugtraq/2001/Mar/182>>.

[Ubuntu2011]

Ubuntu, "USN-1253-1: Linux kernel vulnerabilities", November 2011, <<https://ubuntu.com/security/notices/USN-1253-1>>.

[USCERT2001]

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

[Vixie1995]

Vixie, P., "DNS and BIND Security Issues", 5th Usenix Security Symposium, June 1995, <https://www.usenix.org/legacy/publications/library/proceedings/security95/full_papers/vixie.pdf>.

[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, February 1995.

[Zalewski2001]

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

[Zalewski2002]

Zalewski, M., "Strange Attractors and TCP/IP Sequence Number Analysis - One Year Later (2002)", 2002, <<https://lcamtuf.coredump.cx/newtcp/>>.

[Zalewski2003]

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

Acknowledgements

The authors would like to thank (in alphabetical order) Bernard Aboba, Dave Crocker, Spencer Dawkins, Theo de Raadt, Sara Dickinson, Guillermo Gont, Christian Huitema, Colin Perkins, Vincent Roca, Kris Shrishak, Joe Touch, Brian Trammell, 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 [draft-gont-predictable-numeric-ids-03], on which this document is based.

Section 4.2 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.

Authors' Addresses

Fernando Gont
SI6 Networks
Seguro y Habana
4310 7mo piso
Ciudad Autonoma de Buenos Aires
Argentina
Email: fgont@si6networks.com
URI: <https://www.si6networks.com>

Ivan Arce
Quarkslab
Seguro y Habana
4310 7mo piso
Ciudad Autonoma de Buenos Aires
Argentina
Email: iarce@quarkslab.com
URI: <https://www.quarkslab.com>