Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 4, 2018

F. Gont
SI6 Networks / UTN-FRH
I. Arce
Quarkslab

July 3, 2017

Unfortunate History of Transient Numeric Identifiers
draft-gont-numeric-ids-history-02

Abstract

This document performs an analysis of the security and privacy implications of different types of "numeric identifiers" used in IETF protocols, and tries to categorize them based on their interoperability requirements and the associated failure severity when such requirements are not met. It describes a number of algorithms that have been employed in real implementations to meet such requirements and analyzes their security and privacy properties. Additionally, it provides advice on possible algorithms that could be employed to satisfy the interoperability requirements of each identifier type, while minimizing the security and privacy implications, thus providing guidance to protocol designers and protocol implementers. Finally, it provides recommendations for future protocol specifications regarding the specification of the aforementioned numeric identifiers.

Status of This Memo

Table of Contents

1. Introduction

Network protocols employ a variety of numeric identifiers for different protocol entities, ranging from DNS Transaction IDs (TxIDs) to transport protocol numbers (e.g. TCP ports) or IPv6 Interface Identifiers (IIDs). These identifiers usually have specific properties that must be satisfied such that they do not result in negative interoperability implications (e.g. uniqueness during a specified period of time), and associated failure severities when such properties are not met, ranging from soft to hard failures.

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 leakage that could be exploited for pervasive monitoring [RFC7528]. The root of these issues has been, in many cases, the poor selection of identifiers in such protocols, usually as a result of an insufficient or misleading specification. While it is generally trivial to identify an algorithm that can satisfy the interoperability requirements for a given identifier, there exists practical evidence 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 TCP Initial Sequence Numbers (ISNs) (see e.g. [Morris1985])

- Predictable ephemeral transport protocol numbers (see e.g. [RFC6056] and [Silbersack2005])

- Predictable IPv4 or IPv6 Fragment Identifiers (see e.g. [RFC5722], [RFC6274], and [RFC7739])

- Predictable IPv6 IIDs (see e.g. [RFC7721] and [RFC7707])

- Predictable DNS TxIDs

Recent history indicate that when new protocols are standardized or new protocol implementations are produced, the security and privacy properties of the associated identifiers tend to be overlooked and inappropriate algorithms to generate identifier values are either suggested in the specification or selected by implementers.

This document contains a non-exhaustive timeline of vulnerability disclosures related to some sample transient numeric identifiers and other work that has led to advances in this area, with the goal of illustrating that:

- Vulnerabilities related to how the values for some identifiers are generated and assigned have affected 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.

- Standardization efforts that discuss and provide advice in this area can have a positive effect on protocol specifications and protocol implementations.

Other related documents ([I-D.gont-numeric-ids-generation] and [I-D.gont-numeric-ids-sec-considerations]) provide guidance in this area.

2. Terminology

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. Identifiers are usually defined as a series of bits and represented using integer values. We note that different identifiers may have additional requirements or properties depending on their specific use in a protocol. We use the term "identifier" as a generic term to refer to any data object in a protocol specification that satisfies the identification property stated above.

Failure Severity:
The consequences of a failure to comply with the interoperability requirements of a given identifier. Severity considers the worst potential consequence of a failure, determined by the system damage and/or time lost to repair the failure. In this document we define two types of failure severity: "soft" and "hard".

Hard Failure:
A hard failure is a non-recoverable condition in which a protocol does not operate in the prescribed manner or it operates with excessive degradation of service. For example, an established TCP connection that is aborted due to an error condition constitutes, from the point of view of the transport protocol, a hard failure, since it enters a state from which normal operation cannot be recovered.

Soft Failure:
A soft failure is a recoverable condition in which a protocol does not operate in the prescribed manner but normal operation can be resumed automatically in a short period of time. For example, a simple packet-loss event that is subsequently recovered with a retransmission can be considered a soft failure.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

3. Threat Model

Throughout this document, we assume an attacker does not have physical or logical device to the device(s) being attacked. We assume the attacker can simply send any traffic to the target devices, to e.g. sample identifiers employed by such devices.

4. IPv4/IPv6 Identification

This section presents the timeline of the Identification field both for IPv4 and for IPv6. The reason for presenting both cases in the same section is so that it becomes 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 the IPv6 specifications or implementations.

The IPv4 Identification value 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, some host may be able to simply use unique identifiers independent of destination". The above may be read 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, it misses a discussion of the security and privacy implications of employing predictable. This has resulted in virtually all IP4 implementations generating predictable fragment Identification values by means of a global counter, at least at some point during the lifetime of such implementations.

The IPv6 Identification is specified in [RFC2460]. 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). [RFC2460] 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.

September 1981:
[RFC0791] specifies the interoperability requirements for IPv4 Identification value, but does not specify any requirements in the area of security and privacy.

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] suggests that a global counter be used to generate the IPv6 Identification value.

November 1999:
[Sanfilippo1999] discusses how to leverage predictable IPv4 Identification to uncover the rules of a number of firewalls.

November 1999:
[Bellovin2002] explains how the IPv4 Identification field can be exploited to count the number of systems behind a NAT.

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 is published on the Standards Track, meaning 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 status of the aforementioned IETF Internet Draft 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 on the Informational track, it does not formally update [RFC2460].

June 2016:
[I-D.ietf-6man-rfc2460bis], revision of [RFC2460], removes the suggestion from RFC2460 to employ a global counter for the generation of IPv6 Identification values, but does not specify any

security and privacy requirements for the IPv6 Identification value.

5. 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*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].

In the case of TCP, the interoperability requirements for the ISNs are probably not clearly spelled out as one would expect. Furthermore, the suggestion of employing a global counter in [RFC0793] leads to negative security and privacy implications.

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 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 implications of 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].

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].

February 2012:
[RFC6528], after 27 years of 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.

6. IPv6 Interface Identifiers (IIDs)

IPv6 Interface Identifiers can be generated in multiple ways: SLAAC [RFC4862], DHCPv6 [RFC3315], 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. IPv6 addresses resulting from SLAAC are currently required to employ Modified EUI-64 format identifiers, which essentially embed the underlying link-layer

address of the corresponding network interface. 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 implications of this approach, and thus ongoing work [I-D.ietf-6man-default-iids] at the IETF is in the process of addressing this problem.

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 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 implications on 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] 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 implications, and meaning to eventually provide advice in this area.

7. IANA Considerations

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

8. Security Considerations

The entire document is about the security and privacy implications of transient numeric identifiers.

9. Acknowledgements

The authors would like to thank (in alphabetical order) Dave Crocker, Christian Huitema, and Joe Touch, 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 5 of this document borrows text from [RFC7528], authored by Fernando Gont and Steven Bellovin.

10. References

10.1. Normative References

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <http://www.rfc-editor.org/info/rfc791>.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <http://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, <http://www.rfc-editor.org/info/rfc1323>.

[RFC1884] Hinden, R., Ed. and S. Deering, Ed., "IP Version 6 Addressing Architecture", RFC 1884, DOI 10.17487/RFC1884, December 1995, <http://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, <http://www.rfc-editor.org/info/rfc2073>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2373] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 2373, DOI 10.17487/RFC2373, July 1998, <http://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, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc3041>.

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <http://www.rfc-editor.org/info/rfc3315>.

[RFC3587] Hinden, R., Deering, S., and E. Nordmark, "IPv6 Global Unicast Address Format", RFC 3587, DOI 10.17487/RFC3587, August 2003, <http://www.rfc-editor.org/info/rfc3587>.

[RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, DOI 10.17487/RFC4086, June 2005, <http://www.rfc-editor.org/info/rfc4086>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <http://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, <http://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, <http://www.rfc-editor.org/info/rfc4941>.

[RFC5722] Krishnan, S., "Handling of Overlapping IPv6 Fragments", RFC 5722, DOI 10.17487/RFC5722, December 2009, <http://www.rfc-editor.org/info/rfc5722>.

[RFC6151] Turner, S. and L. Chen, "Updated Security Considerations for the MD5 Message-Digest and the HMAC-MD5 Algorithms", RFC 6151, DOI 10.17487/RFC6151, March 2011, <http://www.rfc-editor.org/info/rfc6151>.

[RFC6528] Gont, F. and S. Bellovin, "Defending against Sequence Number Attacks", RFC 6528, DOI 10.17487/RFC6528, February 2012, <http://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, <http://www.rfc-editor.org/info/rfc7217>.

10.2. 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, <http://www.cert.org/advisories/CA-2001-09.html>.

[CPNI-TCP]
Gont, F., "Security Assessment of the Transmission Control Protocol (TCP)", United Kingdom's Centre for the Protection of National Infrastructure (CPNI) Technical Report, 2009, <http://www.gont.com.ar/papers/tn-03-09-security-assessment-TCP.pdf>.

[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-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-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.

[Fyodor2004]
Fyodor, "Idle scanning and related IP ID games", 2004, <http://www.insecure.org/nmap/idlescan.html>.

[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>, 1995, <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-numeric-ids-generation]
Gont, F. and I. Arce, "On the Generation of Transient Numeric Identifiers", draft-gont-numeric-ids-generation-00 (work in progress), July 2016.

[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-00 (work in progress), June 2016.

[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-00 (work in progress), February 2016.

[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-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-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.

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

[Joncheray1995]
Joncheray, L., "A Simple Active Attack Against TCP", Proc. Fifth Usenix UNIX Security Symposium, 1995.

[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>.

[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, <http://www.rfc-editor.org/info/rfc1035>.

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

[RFC4963] Heffner, J., Mathis, M., and B. Chandler, "IPv4 Reassembly Errors at High Data Rates", RFC 4963, DOI 10.17487/RFC4963, July 2007, <http://www.rfc-editor.org/info/rfc4963>.

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

[RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-Protocol Port Randomization", BCP 156, RFC 6056, DOI 10.17487/RFC6056, January 2011, <http://www.rfc-editor.org/info/rfc6056>.

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

[RFC7528] Higgs, P. and J. Piesing, "A Uniform Resource Name (URN) Namespace for the Hybrid Broadcast Broadband TV (HbbTV) Association", RFC 7528, DOI 10.17487/RFC7528, April 2015, <http://www.rfc-editor.org/info/rfc7528>.

[RFC7707] Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", RFC 7707, DOI 10.17487/RFC7707, March 2016, <http://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, <http://www.rfc-editor.org/info/rfc7721>.

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

[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>.

[Shimomura1995]
Shimomura, T., "Technical details of the attack described by Markoff in NYT", Message posted in USENET's comp.security.misc newsgroup Message-ID: <3g5gkl$5j1@ariel.sdsc.edu>, 1995, <http://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, <http://www.kb.cert.org/vuls/id/498440>.

[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 / UTN-FRH
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

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

Ivan Arce
Quarkslab

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