

Workgroup: Network Working Group
Internet-Draft:
draft-gont-numeric-ids-sec-considerations-07
Updates: [3552](#) (if approved)
Published: 1 February 2022
Intended Status: Best Current Practice
Expires: 5 August 2022
Authors: F. Gont I. Arce
 EdgeUno Quarkslab

Security Considerations for Transient Numeric Identifiers Employed in Network Protocols

Abstract

Poor selection of transient numerical identifiers in protocols such as the TCP/IP suite has historically led to a number of attacks on implementations, ranging from Denial of Service (DoS) to data injection and information leakage that can be exploited by pervasive monitoring. To prevent such flaws in future protocols and implementations, this document updates RFC 3552, requiring future RFCs to contain a vulnerability assessment of their transient numeric identifiers.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 5 August 2022.

Copyright Notice

Copyright (c) 2022 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. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Terminology](#)
- [3. Issues with the Specification of Transient Numeric Identifiers](#)
- [4. Common Flaws in the Generation of Transient Numeric Identifiers](#)
- [5. Vulnerability Assessment Requirements for Transient Numeric Identifiers](#)
- [6. IANA Considerations](#)
- [7. Security Considerations](#)
- [8. Acknowledgements](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

Network protocols employ a variety of transient 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 an associated failure severity when such properties not met.

The TCP/IP protocol suite alone has been subject to variety of attacks on its transient numeric identifiers over the past 30 years or more, with effects ranging from Denial of Service (DoS) or data injection, to information leakage that could be exploited for pervasive monitoring [[RFC7258](#)]. The root of these issues has been, in many cases, the poor selection of 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 for a given identifier, there exists practical evidence [[I-D.irtf-pearl-numeric-ids-history](#)] 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 sequence numbers (see e.g. [[Morris1985](#)], [[Bellare1989](#)], and [[RFC6528](#)])
- *Predictable transport protocol numbers (see e.g. [[Silbersack2005](#)] and [[RFC6056](#)])
- *Predictable IPv4 or IPv6 Fragment Identifiers (see e.g. [[Sanfilippo1998a](#)], [[RFC6274](#)], and [[RFC7739](#)])
- *Predictable IPv6 IIDs (see e.g. [[RFC7217](#)], [[RFC7707](#)], and [[RFC7721](#)])
- *Predictable DNS TxIDs (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 identifiers tend to be overlooked and inappropriate algorithms to generate such identifiers are either suggested in the specification or selected by implementers. As a result, advice in this area is warranted.

We note that the use of cryptographic techniques for confidentiality and authentication may readily mitigate some of the issues arising from predictable transient numeric identifiers. For example, cryptographic authentication can readily mitigate data injection attacks even in the presence of predictable transient numeric identifiers (such as "sequence numbers"). However, use of flawed algorithms (such as global counters) for generating transient numeric identifiers could still result in information leakages even when cryptographic techniques are employed.

[Section 3](#) provides an overview of common flaws in the specification of transient numeric identifiers. [Section 4](#) provides an overview of the implications of predictable transient numeric identifiers. Finally, [Section 5](#) provides key guidelines for protocol designers.

2. Terminology

Transient Numeric Identifier:

A data object in a protocol specification that can be used to definitely distinguish a protocol object (a datagram, network interface, transport protocol endpoint, session, etc) from all other objects of the same type, in a given context. Transient numeric identifiers are usually defined as a series of bits, and represented using integer values. These identifiers are typically dynamically selected, as opposed to statically-assigned numeric identifiers (see e.g. [[IANA-PROT](#)]). We note that different

identifiers may have additional requirements or properties depending on their specific use in a protocol. We use the term "transient numeric identifier" (or simply "numeric identifier" or "identifier" as short forms) as a generic term to refer to any data object in a protocol specification that satisfies the identification property stated above.

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", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

3. Issues with the Specification of Transient Numeric Identifiers

A recent survey of transient numeric identifier usage in protocol specifications and implementations [[I-D.irtf-pearg-numeric-ids-history](#)] revealed that most of the issues discussed in this document arise as a result of one of the following conditions:

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

Both under-specifying and over-specifying identifier contents is hazardous. TCP port numbers and sequence numbers [[RFC0793](#)] and DNS TxID [[RFC1035](#)] were under-specified, leading to implementations that used predictable values and thus were vulnerable to numerous off-path attacks. Over-specification, as for IPv6 Interface Identifiers (IIDs) [[RFC4291](#)] and Fragment Identification values [[RFC2460](#)], leaves implementations unable to respond to security and privacy issues stemming from the mandated algorithm -- IPv6 IIDs need not expose privacy-sensitive link-layer addresses, and predictable Fragment Identifiers invite the same off-path attacks that plague TCP.

Finally, there are protocol implementations that simply fail to comply with existing protocol specifications. That is, appropriate guidance is provided by the protocol specification (whether the core specification or an update to it), but an implementation simply fails to follow such guidance. For example, some popular operating systems (notably Microsoft Windows) still fail to implement transport-protocol port randomization, as specified in [[RFC6056](#)].

Clear specification of the interoperability requirements for the transient numeric identifiers will help identify possible algorithms that could be employed to generate them, and also make evident if such identifiers are being over-specified. A protocol specification will usually also benefit from a vulnerability assessment of the transient numeric identifiers they specify, to prevent the corresponding considerations from being overlooked.

4. Common Flaws in the Generation of Transient Numeric Identifiers

This section briefly notes common flaws associated with the generation of transient numeric identifiers. Such common flaws include, but are not limited to:

- *Employing trivial algorithms (e.g. global counters) that result in predictable identifiers
- *Employing the same identifier across contexts in which constancy is not required
- *Re-using identifiers across different protocols or layers of the protocol stack
- *Initializing counters or timers to constant values, when such initialization is not required
- *Employing the same increment space across different contexts
- *Use of flawed pseudo-random number generators (PRNGs).

Employing trivial algorithms for generating the identifiers means that any node that is able to sample such identifiers can easily predict future identifiers employed by the victim node.

When one identifier is employed across contexts where such constancy is not needed, activity correlation is made made possible. For example, employing an identifier that is constant across networks allows for node tracking across networks.

Re-using identifiers across different layers or protocols ties the security and privacy properties of the protocol re-using the identifier to the security and privacy properties of the original identifier (over which the protocol re-using the identifier may have no control regarding its generation). Besides, when re-using an identifier across protocols from different layers, the goal of isolating the properties of a layer from that of another layer is broken, and the vulnerability assessment may be harder to perform, since the combined system, rather than each protocol in isolation will have to be assessed.

At times, a protocol needs to convey order information (whether sequence, timing, etc.). In many cases, there is no reason for the corresponding counter or timer to be initialized to any specific value e.g. at system bootstrap. Similarly, there may not be a need for the difference between successive counted values to be a predictable.

A node that implements a per-context linear function may share the increment space among different contexts (please see the "Simple Hash-Based Algorithm" in [[I-D.irtf-pearg-numeric-ids-generation](#)]). Sharing the same increment space allows an attacker that can sample identifiers in other context to e.g. learn how many identifiers have been generated between two sampled values.

Finally, some implementations have been found to employ flawed PRNGs (see e.g. [[Klein2007](#)]).

5. Vulnerability Assessment Requirements for Transient Numeric Identifiers

Protocol specifications that employ transient numeric identifiers SHOULD:

1. Clearly specify the interoperability requirements for the aforementioned transient numeric identifiers (e.g., required properties such as uniqueness, along with the failure severity if such properties are not met).
2. Provide a vulnerability assessment of the aforementioned identifiers.

Note: Section 8 and Section 9 of [[I-D.irtf-pearg-numeric-ids-generation](#)] provide a general vulnerability assessment of transient numeric identifiers, along with a vulnerability assessment of common algorithms for generating transient numeric identifiers.

3. Recommend an algorithm for generating the aforementioned transient numeric identifiers that mitigates the vulnerabilities identified in the previous step, such as those discussed in [[I-D.irtf-pearg-numeric-ids-generation](#)].

6. IANA Considerations

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

7. Security Considerations

This document formally updates [[RFC3552](#)] such that a vulnerability assessment of transient numeric identifiers is performed when writing the "Security Considerations" section of future RFCs.

8. Acknowledgements

The authors would like to thank Bernard Aboba, Theo de Raadt, Russ Housley, Benjamin Kaduk, Charlie Kaufman, 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, for providing valuable comments on [[I-D.gont-predictable-numeric-ids](#)] , on which the present document is based.

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

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC3552]

Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <<https://www.rfc-editor.org/info/rfc3552>>.

9.2. Informative References

[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>>.

[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>>.

[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>>.

[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>>.

[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>>.

[Sanfilippo1998a] Sanfilippo, S., "about the ip header id", Post to Bugtraq mailing-list, Mon Dec 14 1998, <<http://seclists.org/bugtraq/1998/Dec/48>>.

[RFC0793]

Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.

[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>>.

[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>>.

[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>>.
- [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>>.
- [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>>.
- [Silbersack2005] Silbersack, M.J., "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>>.
- [I-D.gont-predictable-numeric-ids] 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://www.ietf.org/archive/id/draft-gont-predictable-numeric-ids-03.txt>>.
- [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>>.
- [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>>.
- [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>>.
- [Klein2007] Klein, A., "OpenBSD DNS Cache Poisoning and Multiple O/S Predictable IP ID Vulnerability", 2007, <https://dl.packetstormsecurity.net/papers/attack/OpenBSD_DNS_Cache_Poisoning_and_Multiple_OS_Predictable_IP_ID_Vulnerability.pdf>.
- [Schuba1993] Schuba, C., "ADDRESSING WEAKNESSES IN THE DOMAIN NAME SYSTEM PROTOCOL", 1993, <<http://ftp.cerias.purdue.edu/pub/papers/christoph-schuba/schuba-DNS-msthesis.pdf>>.
- [I-D.irtf-pearg-numeric-ids-history] Gont, F. and I. Arce, "Unfortunate History of Transient Numeric Identifiers", Work in Progress, Internet-Draft, draft-irtf-pearg-

numeric-ids-history-09, 27 January 2022, <<https://www.ietf.org/archive/id/draft-irtf-pearg-numeric-ids-history-09.txt>>.

[I-D.irtf-pearg-numeric-ids-generation] Gont, F. and I. Arce, "On the Generation of Transient Numeric Identifiers", Work in Progress, Internet-Draft, draft-irtf-pearg-numeric-ids-generation-07, 2 February 2021, <<https://www.ietf.org/archive/id/draft-irtf-pearg-numeric-ids-generation-07.txt>>.

[IANA-PROT] IANA, "Protocol Registries", <<https://www.iana.org/protocols>>.

Authors' Addresses

Fernando Gont
EdgeUno
Segurola y Habana 4310 7mo piso
Ciudad Autonoma de Buenos Aires
Buenos Aires
Argentina

Email: fernando.gont@edgeuno.com
URI: <https://www.edgeuno.com>

Ivan Arce
Quarkslab
Segurola y Habana 4310 7mo piso
Ciudad Autonoma de Buenos Aires
Buenos Aires
Argentina

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