

TCP Maintenance and Minor
Extensions (tcpm)
Internet-Draft
Intended status: BCP
Expires: May 1, 2009

F. Gont
Consultant
October 28, 2008

On the generation of TCP timestamps
draft-gont-tcpm-tcp-timestamps-00.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79. This document may not be modified, and derivative works of it may not be created.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on May 1, 2009.

Abstract

This document describes an algorithm for selecting the timestamps (TS value) used for TCP connections that use the TCP timestamp option, such that the resulting timestamps are monotonically-increasing across connections that involve the same four-tuple {local IP address, local TCP port, remote IP address, remote TCP port}. The properties of the algorithm are such the possibility of an attacker guessing the exact value is reduced. Additionally, it describes an algorithm for processing incoming SYN segments to allow a higher connection establishment rates to any TCP end-point.

Table of Contents

- 1. Introduction 3
- 2. Proposed algorithm 3
- 3. Improved processing of incoming connection requests 4
- 4. Security Considerations 6
- 5. IANA Considerations 6
- 6. Acknowledgements 6
- 7. References 7
 - 7.1. Normative References 7
 - 7.2. Informative References 7
- Author's Address 7
- Intellectual Property and Copyright Statements 8

1. Introduction

The Timestamps option, specified in RFC 1323 [RFC1323], allows a TCP to include a timestamp value in its segments, that can be used used to perform two functions: Round-Trip Time Measurement (RTTM), and Protect Against Wrapped Sequences (PAWS).

For the purpose of PAWS, the timestamps sent on a connection are required to be monotonically increasing. While there is no requirement that timestamps are monotonically increasing across TCP connections, the generation of timestamps such that they are monotonically increasing across connections between the same two endpoints allows the use of timestamps for improving the handling of SYN segments that are received while the corresponding four-tuple is in the TIME-WAIT state. That is, the timestamp option could be used to perform heuristics to determine whether to allow the creation of a new incarnation of a connection that is in the TIME-WAIT state.

This use of TCP timestamps is simply an extrapolation of the use of ISNs for the same purpose, as allowed by RFC 1122 [RFC0793] itself, and has been incorporated in a number of TCP implementations, such as that included in the Linux kernel. [Linux]

In order to avoid the security implications of predictable timestamps, the proposed algorithm generates timestamps such such that the possibility of an attacker guessing the exact value is reduced.

Section 2 proposes the aforementioned algorithm for generating TCP timestamps. Section 3 describes an improved processing of incoming connection requests, that may allow higher connection-establishment rates to any TCP end-point.

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].

2. Proposed algorithm

It is RECOMMENDED that timestamps are generated with a similar algorithm to that introduced by RFC 1948 [RFC1948] for the generation of Initial Sequence Numbers (ISNs). That is,

```
timestamp = T() + F(localhost, localport, remotehost, remoteport,  
secret_key)
```

where the result of T() is a global system clock that complies with

the requirements of Section 4.2.2 of RFC 1323 [RFC1323], and $F()$ is a function that should not be computable from the outside. Therefore, we suggest $F()$ to be a cryptographic hash function of the connection-id and some secret data.

$F()$ provides an offset that will be the same for all incarnations of a connection between the same two endpoints, while $T()$ provides the monotonically increasing values that are needed for PAWS.

3. Improved processing of incoming connection requests

In a number of scenarios a socket pair may need to be reused while the corresponding four-tuple is still in the TIME-WAIT state in a remote TCP peer. For example, a client accessing some service on a host may try to create a new incarnation of a previous connection, while the corresponding four-tuple is still in the TIME-WAIT state at the remote TCP peer (the server). This may happen if the ephemeral port numbers are being reused too quickly, either because of a bad policy of selection of ephemeral ports, or simply because of a high connection rate to the corresponding service. In such scenarios, the establishment of new connections that reuse a four-tuple that is in the TIME-WAIT state would fail. In order to avoid this problem, RFC 1122 [RFC1122] (in Section 4.2.2.13) states that when a connection request is received with a four-tuple that is in the TIME-WAIT state, the connection request could be accepted if the sequence number of the incoming SYN segment is greater than the last sequence number seen on the previous incarnation of the connection (for that direction of the data transfer).

This requirement aims at avoiding the sequence number space of the new and old incarnations of the connection to overlap, thus avoiding old segments from the previous incarnation of the connection to be accepted as valid by the new connection.

The following paragraphs summarize the processing of SYN segments received for connections in the TIME-WAIT state. Both the ISN (Initial Sequence Number) and the timestamp option (if present) of the incoming SYN segment are included in the heuristics performed for allowing a high connection-establishment rate.

Processing of SYN segments received for connections in the TIME-WAIT state should occur as follows:

- o If the previous incarnation of the connection used timestamps, then,

- * If the incoming SYN segment contains a timestamp option, and the timestamp is greater than the last timestamp seen on the previous incarnation of the connection, honor the connection request (creating a connection in the SYN-RECEIVED state).
 - * If the incoming SYN segment contains a timestamp option, and the timestamp is equal to the last timestamp seen on the previous incarnation of the connection (for that direction of the data transfer), and the sequence number of the incoming SYN segment is larger than the last sequence number seen on the previous incarnation of the connection (for that direction of the data transfer), then honor the connection request.
 - * If the incoming SYN segment does not contain a timestamp option, but the Sequence Number of the incoming SYN segment is larger than the last sequence number seen on the previous incarnation of the connection (for the same direction of the data transfer), honor the connection request.
 - * Otherwise, silently drop the incoming SYN segment, thus leaving the previous connection in the TIME-WAIT state.
- o If the previous incarnation of the connection did not use timestamps, then,
 - * If the incoming connection request contains a timestamp option, and timestamps will be enabled for the new incarnation of the connection (i.e., the SYN/ACK segment will contain a timestamp option), honor the incoming connection request.
 - * If the incoming connection request does not contain a timestamp option, but the Sequence Number of the incoming SYN segment is larger than the last sequence number seen on the previous incarnation of the connection for the same direction of the data transfer, then honor the incoming connection request (even if the sequence number of the incoming SYN segment falls within the receive window of the previous incarnation of the connection).

Many implementations do not include the TCP timestamp option when performing the above heuristics, thus imposing stricter constraints on the generation of Initial Sequence Numbers, the average data transfer rate of the connections, and the amount of data transferred with them. RFC 793 [RFC0793] states that the ISN generator should be incremented roughly once every four microseconds (i.e., roughly 250000 times per second). As a result, any connection that transfers more than 250000 bytes of data at more than 250 KB/s could lead to scenarios in which the last sequence number seen on a connection that

moves into the TIME-WAIT state is still greater than the sequence number of an incoming SYN segment that aims at creating a new incarnation of the same connection. In those scenarios, the 4.4BSD heuristics would fail, and therefore the connection request would usually time out. By including the TCP timestamp option in the heuristics described above, all these constraints are greatly relaxed.

It is clear that the use of TCP timestamps for the heuristics described above depends on the timestamps to be monotonically increasing across connections between the same two TCP endpoints. Therefore, we strongly advice to generate timestamps as described in Section 2.

4. Security Considerations

This document describes an algorithm that can be used to obfuscate the timestamp value used for new connections, such that the possibility of an attacker guessing the exact value is reduced.

Some implementations are known to maintain a global timestamp clock, which is used for all connections. This is undesirable, as an attacker that can establish a connection with a host would learn the timestamp used for all the other connections maintained by that host, which could be useful for performing any attacks that require the attacker to forge TCP segments. Some implementations are known to initialize their global timestamp clock to zero when the system is bootstrapped. This is undesirable, as the timestamp clock would disclose the system uptime.

The algorithm discussed in this document for generating the TCP timestamps avoids these problems by generating timestamps as monotonically-increasing function with a per-connection-id random offset. [CPNI-TCP]

5. IANA Considerations

This document has no actions for IANA.

6. Acknowledgements

Yet to be added

7. References

7.1. Normative References

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.
- [RFC1323] Jacobson, V., Braden, B., and D. Borman, "TCP Extensions for High Performance", RFC 1323, May 1992.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

- [CPNI-TCP] CPNI, "Security Assessment of the Transmission Control Protocol (TCP)", (to be published) .
- [Linux] The Linux Project, "<http://www.kernel.org>".
- [RFC1948] Bellovin, S., "Defending Against Sequence Number Attacks", RFC 1948, May 1996.

Author's Address

Fernando Gont
Consultant

Email: fernando@gont.com.ar
URI: <http://www.gont.com.ar>

Full Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

