

TCP Maintenance and Minor
Extensions (tcpm)
Internet-Draft
Intended status: BCP
Expires: June 23, 2011

F. Gont
UK CPNI
December 20, 2010

Reducing the TIME-WAIT state using TCP timestamps
draft-ietf-tcpm-tcp-timestamps-03.txt

Abstract

This document describes an algorithm for processing incoming SYN segments that allows higher connection-establishment rates between any two TCP endpoints when a TCP timestamps option is present in the incoming SYN segment. This document only modifies processing of SYN segments received for connections in the TIME-WAIT state; processing in all other states is unchanged.

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 <http://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 June 23, 2011.

Copyright Notice

Copyright (c) 2010 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 (<http://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 Simplified BSD License text as described in [Section 4.e](#) of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	3
2.	Improved processing of incoming connection requests	3
3.	Interaction with various timestamps generation algorithms	6
4.	Interaction with various ISN generation algorithms	7
5.	Security Considerations	8
6.	IANA Considerations	8
7.	Acknowledgements	8
8.	References	8
8.1.	Normative References	8
8.2.	Informative References	9
Appendix A.	Behavior of the proposed mechanism in specific scenarios	9
A.1.	Connection request after system reboot	10
Appendix B.	Changes from previous versions of the draft (to be removed by the RFC Editor before publishing this document as an RFC)	10
B.1.	Changes from draft-ietf-tcpm-tcp-timestamps-02	10
B.2.	Changes from draft-ietf-tcpm-tcp-timestamps-01	10
B.3.	Changes from draft-ietf-tcpm-tcp-timestamps-00	10
B.4.	Changes from draft-gont-tcpm-tcp-timestamps-04	10
B.5.	Changes from draft-gont-tcpm-tcp-timestamps-03	11
B.6.	Changes from draft-gont-tcpm-tcp-timestamps-02	11
B.7.	Changes from draft-gont-tcpm-tcp-timestamps-01	11
B.8.	Changes from draft-gont-tcpm-tcp-timestamps-00	11
	Author's Address	11

Gont

Expires June 23, 2011

[Page 2]

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 to perform two functions: Round-Trip Time Measurement (RTTM), and Protection 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 Initial Sequence Numbers (ISNs) for the same purpose, as allowed by [RFC 1122](#) [[RFC1122](#)], and it has been incorporated in a number of TCP implementations, such as that included in the Linux kernel [[Linux](#)].

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. 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. This problem is discussed in detail in [[INFOCOM-99](#)].

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 same policy may be extrapolated to TCP timestamps. That is, 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 timestamp of the incoming SYN segment is greater than the last timestamp seen on the previous incarnation of the connection (for that direction of the data transfer).

The following paragraphs summarize the processing of SYN segments received for connections in the TIME-WAIT state. The processing of SYN segments received for connections in all other states is unchanged. Both the ISN (Initial Sequence Number) and the timestamps 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 TCP timestamps would be enabled for the new incarnation of the connection, and the timestamp contained in the incoming SYN segment is greater than the last timestamp seen on the previous incarnation of the connection (for that direction of the data transfer), honour the connection request (creating a connection in the SYN-RECEIVED state).
 - * If TCP timestamps would be enabled for the new incarnation of the connection, the timestamp contained in the incoming SYN segment 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 greater than the last sequence number seen on the previous incarnation of the connection (for that direction of the data transfer), honour the connection request (creating a connection in the SYN-RECEIVED state).
 - * If TCP timestamps would not be enabled for the new incarnation of the connection, but the Sequence Number of the incoming SYN segment is greater than the last sequence number seen on the previous incarnation of the connection (for the same direction

- of the data transfer), honour the connection request (creating a connection in the SYN-RECEIVED state).
- * Otherwise, silently drop the incoming SYN segment, thus leaving the previous incarnation of the connection in the TIME-WAIT state.
- o If the previous incarnation of the connection did not use timestamps, then,
 - * If TCP timestamps would be enabled for the new incarnation of the connection, honour the incoming connection request (creating a connection in the SYN-RECEIVED state).
 - * If TCP timestamps would not be enabled for the new incarnation of the connection, but the Sequence Number of the incoming SYN segment is greater than the last sequence number seen on the previous incarnation of the connection (for the same direction of the data transfer), honour the incoming connection request (creating a connection in the SYN-RECEIVED state).
 - * Otherwise, silently drop the incoming SYN segment, thus leaving the previous incarnation of the connection in the TIME-WAIT state.

Note:

In the above explanation, the phrase "TCP timestamps would be enabled for the new incarnation for the connection" means that the incoming SYN segment contains a TCP Timestamps option (i.e., the client has enabled TCP timestamps), and that the SYN/ACK segment that would be sent in response to it would also contain a Timestamps option (i.e., the server has enabled TCP timestamps). In such a scenario, TCP timestamps would be enabled for the new incarnation of the connection.

The "last sequence number seen on the previous incarnation of the connection (for the same direction of the data transfer)" refers to the last sequence number used by the previous incarnation of the connection (for the same direction of the data transfer), and not to the last value seen in the Sequence Number field of the corresponding segments. That is, it refers to the sequence number corresponding to the FIN flag of the previous incarnation of the connection, for that direction of the data transfer.

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 kilobytes/second 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 benefit from timestamps that are monotonically increasing across connections between the same two TCP endpoints.

Note:

The upcoming revision of [RFC 1323](#), [[I-D.ietf-tcpm-1323bis](#)], recommends the selection of timestamps such that they are monotonically-increasing across connections. An example of such a Timestamps generation scheme can be found in [[I-D.gont-timestamps-generation](#)].

3. Interaction with various timestamps generation algorithms

The algorithm proposed in [Section 2](#) clearly benefits of timestamps that are monotonically-increasing across connections to the same endpoint. In particular, generation of timestamps such that they are monotonically-increasing timestamps are important for TCPs that perform the active open, as those are the timestamps that will be used for the proposed algorithm.

While monotonically-increasing timestamps ensure that the proposed algorithm will be able to reduce the TIME-WAIT state of a previous incarnation of a connection, implementation of the algorithm does not imply by itself a requirement on the timestamps generation algorithm of other TCPs.

In the worst-case scenario, an incoming SYN corresponding to a new incarnation of a connection in the TIME-WAIT contains a timestamp that is smaller than the last timestamp seen on the previous incarnation of the connection, the heuristics fail, and the result is no worse than the current state-of-affairs. That is, the SYN segment is ignored (as specified in [[RFC1337](#)]), and thus the connection request times out, or is accepted after future retransmissions of the

SYN.

Some stacks may implement timestamps generation algorithms that do not lead to monotonically-increasing timestamps across connections with the same remote endpoint. An example of such algorithms is the one described in [\[RFC4987\]](#) and [\[Opperman\]](#), that allows the implementation of extended TCP SYN cookies.

Note:

It should be noted that the "extended TCP SYN cookies" could co-exist with an algorithm for generating timestamps such that they are monotonically-increasing. Monotonically increasing timestamps could be generated for TCPs that perform the active open, while timestamps for TCPs that perform the passive open could be generated according to [\[Opperman\]](#).

Some stacks (notably OpenBSD) implement timestamps randomization algorithms which do not result in monotonically-increasing ISNs across connections. As noted in [\[Silbersack\]](#), such randomization schemes may prevent the mechanism proposed in this document from recycling connections that are in the TIME-WAIT state. However, as noted earlier in this section, in the worst-case scenario the heuristics fail, and the result is no worse than the current state-of-affairs.

4. Interaction with various ISN generation algorithms

[\[RFC0793\]](#) suggests that the ISNs of TCP connections be generated from a global timer, such that they are monotonically-increasing across connections. However, this ISN-generation scheme leads to predictable ISNs, which have well-known security implications [\[CPNI-TCP\]](#). [\[RFC1948\]](#) proposes an alternative ISN-generation scheme which results in monotonically-increasing ISNs across connections that are not easily-predictable by an off-path attacker.

Some stacks (notably OpenBSD) implement ISN randomization algorithms which do not result in monotonically-increasing ISNs across connections. As noted in [\[Silbersack\]](#), such ISN randomization schemes break the BSD improved handling of SYN segments received for connections that are in the TIME-WAIT state.

An implementation of the mechanism proposed in this document would enable recycling of the TIME-WAIT state even in the presence of ISNs that are not monotonically-increasing across connections, except when the timestamp contained in the incoming SYN is equal to the last timestamp seen on the connection in the TIME-WAIT state (for that direction of the data transfer).

5. Security Considerations

While the algorithm described in this document for processing incoming SYN segments would benefit from TCP timestamps that are monotonically-increasing across connections, this document does not propose any specific algorithm for generating timestamps, nor does it require monotonically-increasing timestamps across connections. [CPNI-TCP] contains a detailed discussion of the security implications of TCP timestamps and of different Timestamps generation algorithms.

6. IANA Considerations

This document has no actions for IANA.

7. Acknowledgements

The author of this document would like to thank (in alphabetical order) Mark Allman, Francis Dupont, Wesley Eddy, Lars Eggert, Alfred Hoenes, John Heffner, Christian Huitema, Eric Rescorla, Joe Touch, and Alexander Zimmermann for providing valuable feedback on an earlier version of this document.

Additionally, the author would like to thank David Borman for a fruitful discussion on TCP timestamps at IETF 73.

Finally, the author would like to thank the United Kingdom's Centre for the Protection of National Infrastructure (UK CPNI) for their continued support.

8. References

8.1. 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.

8.2. Informative References

- [CPNI-TCP]
CPNI, "Security Assessment of the Transmission Control Protocol (TCP)", <http://www.cpni.gov.uk/Docs/tn-03-09-security-assessment-TCP.pdf>, 2009.
- [I-D.gont-timestamps-generation]
Gont, F. and A. Oppermann, "On the generation of TCP timestamps", [draft-gont-timestamps-generation-00](#) (work in progress), June 2010.
- [I-D.ietf-tcpm-1323bis]
Borman, D., Braden, R., and V. Jacobson, "TCP Extensions for High Performance", [draft-ietf-tcpm-1323bis-01](#) (work in progress), March 2009.
- [INFOCOM-99]
Faber, T., Touch, J., and W. Yue, "The TIME-WAIT state in TCP and Its Effect on Busy Servers", Proc. IEEE Infocom, 1999, pp. 1573-1583 .
- [Linux] The Linux Project, "<http://www.kernel.org>".
- [Opperman]
Oppermann, A., "FYI: Extended TCP syncookies in FreeBSD-current", Post to the tcpm mailing-list. Available at: <http://www.ietf.org/mail-archive/web/tcpm/current/msg02251.html>, 2006.
- [RFC1337] Braden, B., "TIME-WAIT Assassination Hazards in TCP", [RFC 1337](#), May 1992.
- [RFC1948] Bellovin, S., "Defending Against Sequence Number Attacks", [RFC 1948](#), May 1996.
- [RFC4987] Eddy, W., "TCP SYN Flooding Attacks and Common Mitigations", [RFC 4987](#), August 2007.
- [Silbersack]
Silbersack, M., "Improving TCP/IP security through randomization without sacrificing interoperability", EuroBSDCon 2005 Conference .

Appendix A. Behavior of the proposed mechanism in specific scenarios

A.1. Connection request after system reboot

This section clarifies how this algorithm would operate in case a computer reboots, keeps the same IP address, loses memory of the previous timestamps, and then tries to reestablish a previous connection.

Firstly, as specified in [\[RFC0793\]](#), hosts must not establish new connections for a period of $2 \times \text{MSL}$ (Maximum Segment Lifetime) after they boot (this is the "quiet time" concept). As a result, spec-wise, this scenario should never occur.

If a host does not comply with the "quiet time concept", a connection request might be sent to a remote host while there is a previous incarnation of the same connection in the TIME-WAIT state at the remote host. In such a scenario, as a result of having lost memory of previous time stamps, the resulting timestamps might not be monotonically-increasing, and hence the proposed algorithm might be unable to recycle the previous incarnation of the connection that is in the TIME-WAIT state. This case corresponds to the current state-of-affairs without the algorithm proposed in this document.

Appendix B. Changes from previous versions of the draft (to be removed by the RFC Editor before publishing this document as an RFC)

B.1. Changes from [draft-ietf-tcpm-tcp-timestamps-02](#)

- o Addresses COMMENTS received during IESG review, and maybe Tim Polk's DISCUSS.

B.2. Changes from [draft-ietf-tcpm-tcp-timestamps-01](#)

- o Addresses AD-review comments by Lars Eggert.

B.3. Changes from [draft-ietf-tcpm-tcp-timestamps-00](#)

- o Addresses WG Last call comments received from Wesley Eddy, John Heffner and Joe Touch.
- o Minor editorial fix (reported by Wes Eddy).

B.4. Changes from [draft-gont-tcpm-tcp-timestamps-04](#)

- o Draft resubmitted as [draft-ietf](#).

B.5. Changes from [draft-gont-tcpm-tcp-timestamps-03](#)

- o Changed the document title
- o Removed all the text related to the algorithm earlier proposed for timestamps generation.
- o Addresses comments received from Alexander Zimmermann, Christian Huitema, Joe Touch, and others.

B.6. Changes from [draft-gont-tcpm-tcp-timestamps-02](#)

- o Minor edits (the I-D was just about to expire, so it was resubmitted with almost no changes).

B.7. Changes from [draft-gont-tcpm-tcp-timestamps-01](#)

- o Version -01 of the draft had expired, and hence the I-D is resubmitted to make it available again (no changes).

B.8. Changes from [draft-gont-tcpm-tcp-timestamps-00](#)

- o Fixed author's affiliation.
- o Addressed feedback submitted by Alfred Hoenes (see: <http://www.ietf.org/mail-archive/web/tcpm/current/msg04281.html>), plus nits sent by Alfred off-list.

Author's Address

Fernando Gont
UK Centre for the Protection of National Infrastructure

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

