TCP Adaptive User TimeOut (AUTO) Option
draft-gont-tcpm-tcp-auto-option-00.txt

Status of this memo

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/1id-abstracts.html.

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

This Internet-Draft will expire on November 19, 2004.

Copyright Notice

Abstract

The original TCP specification (RFC 793) defines a "USER TIMEOUT"
parameter that sets the policy as to when a user connection should be
aborted. However, TCP provides no means of letting users suggest an
abort policy to a remote peer dynamically. Even though a fixed policy
may work well in many cases, there are a number of scenarios where a
fixed USER TIMEOUT value may be inappropriate, and some means of
setting the abort policy dynamically may be necessary for TCP to be
used effectively in such scenarios. This document defines a new TCP
option, which lets a TCP peer suggest a USER TIMEOUT value to a
remote TCP during the connection-establishment phase, and modify it

during the life of a connection, thus adapting TCP's connection-abort
policy as necessary.

1. Introduction

The original TCP specification [1] defines a USER TIMEOUT parameter,
which sets the policy as to when a connection should be aborted. This
parameter is usually set on a per-system basis, and there is no way
for a TCP to suggest a value of USER TIMEOUT to be used for a
connection by a remote peer.

Even though having such a fixed policy may work well in many cases,
there are scenarios in which the default USER TIMEOUT may be
inappropriate. For example, a mobile host connected to a network by
means of a wireless link may experience transient periods of
disconnection that may be longer than the USER TIMEOUT selected by
the remote peer. Another possible scenario is the development of
high levels of congestion during the life of a connection.
In such cases, valid connections may be aborted due to an incorrect
abort policy.

This document defines a new TCP option that lets TCP implementations
suggest a USER TIMEOUT value during the connection-establishment
phase, and modify it during the life of a connection, thus adapting
TCP's connection-abort policy as necessary.

2. Conventions

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 [2].

3. Option Format

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Kind = X  |   Length = 4  |G|        User Timeout        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

      Note that one tick mark represents one bit position

          Figure 1: Adaptive User Timeout Option Format

   Each field is to be interpreted as follows:

   Kind: 8 bits
      This is the "Kind" field as specified in [1]. The "X" in Figure 1
      is an option number to be assigned by IANA upon publication of
      this document (see Section 7)

   Length: 8 bits
      This is the "Length" field as specified in [1]. Its value is 4
      (the option length).

   G: 1 bit
      This is the "Granularity" bit. It indicates the granularity of the
      "User Timeout" field. When set, the time interval in the "User
      Timeout" field MUST be interpreted as being specified in minutes.
      Otherwise, the time interval in the "User Timeout" field MUST be
      interpreted as being specified in seconds.

   User Timeout: 15 bits
      This field, together with the Granularity bit, specifies the USER
      TIMEOUT suggested by the remote peer for this connection. It MUST
      be interpreted as a 15-bit unsigned integer. The units of this
      field are specified by the "G" bit.

3. Operation

   TCP implementations supporting the Adaptive User TimeOut (AUTO)
   Option MUST set this option during the connection-establishment phase
   (in segments with the SYN control bit set) to indicate the suggested
   USER TIMEOUT value to be used for the connection.

   A TCP MAY also use this option during the life of a connection, to
   suggest an a new value for the USER TIMEOUT parameter, thus adapting
   it to the current network conditions. For example, this option could
   be set by a TCP peer that is notified of congestion by means of
   ECN [5].

The setting of this option means "I suggest we use a USER TIMEOUT
of X". The value of "X" may be larger or smaller than the default
USER TIMEOUT (see Section 4).

Hosts SHOULD impose upper and lower limits on the USER TIMEOUT. A
discussion of these limits can be found in Section 5.

Each TCP will adopt a USER TIMEOUT as defined by equation (1):

USER_TIMEOUT = min(ULimit, max(localAUTO, remoteAUTO, LLimit))   (1)

USER_TIMEOUT:
    USER TIMEOUT value to be adopted by the local TCP for this
    connection.

Ulimit:
    The upper limit imposed by this host for the USER TIMEOUT.

Llimit:
    The lower limit imposed by this host for the USER TIMEOUT.

localAUTO:
    The "USER TIMEOUT" value suggested by the local TCP by means of
    the AUTO Option.

remoteAUTO:
    The "USER TIMEOUT" value suggested by the remote TCP peer by means
    of the AUTO Option.

The adopted USER TIMEOUT SHOULD be used only for connections that are
in one of the synchronized states (ESTABLISHED, FIN-WAIT-1,
FIN-WAIT-2, CLOSE-WAIT, CLOSING, LAST-ACK or TIME-WAIT).

Note that the USER TIMEOUT is not negotiated in any way. Each peer
just "suggests" what USER TIMEOUT should be adopted for the
connection. As can be inferred from the equation above, each peer may
end up adopting a different timeout value.

4. Range of valid values

The User Timeout Option allows a TCP peer to suggest USER TIMEOUT
values ranging, in principle, from 0 seconds to about 22.76 days.
However, implementations SHOULD impose limits on the USER TIMEOUT
values actually adopted. A discussion of these limits can be found
in Section 5.

5. System limits on the USER TIMEOUT

Implementations SHOULD impose an upper limit (Ulimit) and a lower
limit (Llimit) on the value of the USER TIMEOUT. These limits could,
for example, be set on a per-host or per-user basis.

Furthermore, these limits need not be fixed. For example, they MAY
be a function of the system resources that are available when the
USER TIMEOUT is to be selected for a connection.

The Host Requirements RFC [3] does not impose any limits for the USER
TIMEOUT. However, a time interval of at least 100 seconds is
RECOMMENDED. Thus, the lower limit (LLimit) should be set to at least
100 seconds. As for the upper limit (ULimit), note that setting it to
low values may reduce the functionality of the AUTO Option.

6. Interoperability issues

6.1 Firewalls

Stateful firewalls are known to reset connections after some fixed
period of inactivity is detected. In case there is such a firewall
between the TCP peers, then, regardless of the use of the AUTO
Option, connections may be lost due to the firewall policy.

6.2 TCP Keep-alive mechanism

In case a TCP peer enables the TCP Keep-alive mechanism for a
connection that is using the AUTO Option, then the Keep-alive timer
MUST be set to a value larger than that of the adopted USER
TIMEOUT (specified by Equation 1).

7. IANA Considerations

This section is to be interpreted according to [4].

This document does not define any new namespaces. It uses an 8-bit
TCP option number maintained by IANA at
http://www.iana.org/assignments/tcp-parameters.

8. Security Implications

Use of the AUTO Option implies that the adopted USER TIMEOUT be
larger than the default USER TIMEOUT. This could cause a host to
maintain state for a connection for a longer period of time than if
the default USER TIMEOUT were used. An attacker could try to exhaust
resources on the target host by establishing lots of connections
and aborting them without signalling this to the attacked host's TCP.
However, it must be noted that the same type of attack could be
performed even if the default "USER TIMEOUT" is being used, since TCP
requires no message exchange in order to keep a connection open.
In any case, the system limits discussed in Section 5 would serve
as a counter-measure against attackers trying to exploit the AUTO
option for this type of attack.

9. Author's address

Fernando Gont
Evaristo Carriego 2644
1706, Haedo
Provincia de Buenos Aires
ARGENTINA

Phone: +54 011 4650 8472
E-Mail: fernando@gont.com.ar

10. Acknowledgements

The author wishes to thank Michael Kerrisk for contributing many
valuable comments.

11. References

11.1 Normative References

[1] Postel, J., "Transmission Control Protocol", STD 7, RFC 793,
    September 1981.

[2] Bradner, S., "Key words for use in RFCs to Indicate Requirement
    Levels", BCP 14, RFC 2119, March 1997.

[3] Braden, R., "Requirements for Internet Hosts - Communication
    Layers", STD 3, RFC 1122, October 1989.

[4] Narten, T., Alvestrand, H., "Guidelines for Writing an IANA
    Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

[11.2](#) Informative References

[5] Ramakrishnan, K., Floyd, S., Black, D., "The Addition of Explicit
    Network Notification (ECN) to IP", [RFC 3168](#), September 2001.

Intellectual Property Statement

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.

Disclaimer of Validity

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 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.

Copyright Statement