

TCPM Working Group
Internet Draft
Intended status: Proposed Standard
Expires: August 2013

J. Touch
USC/ISI
February 25, 2013

Shared Use of Experimental TCP Options
draft-ietf-tcpm-experimental-options-04.txt

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), 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/1id-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 August 25, 2013.

Copyright Notice

Copyright (c) 2013 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.

Abstract

This document describes how the experimental TCP option codepoints can concurrently support multiple TCP extensions, even within the same connection. It uses a new IANA TCP experiment identifier, and is also robust to experiments that are not registered and those that do not use this sharing mechanism. It is recommended for all new TCP options that use these codepoints.

Table of Contents

1.	Introduction.....	2
2.	Conventions used in this document.....	4
3.	TCP Experimental Option Structure.....	4
3.1.	Selecting an ExID.....	5
3.2.	Impact on TCP Option Processing.....	6
4.	Reducing the Impact of False Positives.....	6
5.	Migration to Assigned Options.....	7
6.	Security Considerations.....	7
7.	IANA Considerations.....	7
8.	References.....	8
8.1.	Normative References.....	8
8.2.	Informative References.....	8
9.	Acknowledgments.....	9

[1. Introduction](#)

TCP includes options to enable new protocol capabilities that can be activated only where needed and supported [[RFC793](#)]. The space for identifying such options is small - 256 values, of which 30 are assigned at the time this document was published [[IANA](#)]. Two of these codepoints are allocated to support experiments (253, 254) [[RFC4727](#)]. These values are intended for testing purposes or anytime an assigned codepoint is either not warranted or available, e.g., based on the maturity status of the defined capability (i.e., Experimental or Informational, rather than Standards Track).

The term "experimental TCP options" refers here to options that use the TCP experimental option codepoints [[RFC4727](#)]. Such experiments can be described in any type of RFC - Experimental, Informational, etc., and are intended to be used both in controlled environments and in are allowed in public deployments (when not enabled as

Touch

Expires August 25, 2013

[Page 2]

default) [[RFC3692](#)]. Nothing prohibits deploying multiple experiments in the same environment - controlled or public. Further, some protocols are specified in Experimental or Informational RFCs, which either include parameters or design choices not yet understood or which might not be widely deployed [[RFC2026](#)]. TCP options in such RFCs are typically not eligible for assigned TCP option codepoints [[RFC2780](#)], and so there is a need to share use of the experimental option codepoints.

There is currently no mechanism to support shared use of the TCP experimental option codepoints, either by different experiments on different connections, or for more than two experimental options in the same connection. Experimental options 253 and 254 are already deployed in operational code to support an early version of TCP authentication. Option 253 is also documented for the experimental TCP Cookie Transaction option [[RFC6013](#)]. This shared use results in collisions in which a single codepoint can appear multiple times in a single TCP segment and for which each use is ambiguous.

Other codepoints have been used without assignment (known as "squatting"), notably 31-32 (TCP cookie transactions, as originally distributed and in its API doc) and 76-78 (tcpcrypt) [[Bi11](#)][[Si11](#)]. Commercial products reportedly also use unassigned options 33, 69-70, and 76-78 as well. Even though these uses are unauthorized, they currently impact legitimate assignees.

Both such misuses (squatting on both experimental and assigned codepoints) are expected to continue, but there are several approaches which can alleviate the impact on cooperating protocol designers. One proposal relaxes the requirements for assignment of TCP options, allowing them to be assigned more readily for protocols that have not been standardized through the IETF process [[RFC5226](#)]. Another proposal assigns a larger pool to the TCP experiment option codepoints and manages their sharing through IANA coordination [[Ed11](#)].

The approach proposed in this document does not require additional TCP option codepoints, and is robust to those who choose either not to support it or not to register their experiments. The solution adds a field to the structure of the experimental TCP option. This field is populated with an "experiment identifier" (ExID) defined as part of a specific option experiment. The ExID helps reduce the probability of a collision of independent experimental uses of the same option codepoint, both for those who follow this document (using registered ExIDs) and those who do not (squatters who either ignore this extension or do not register their ExIDs).

The solution proposed in this document is recommended for all new protocols that use TCP experimental option codepoints. The techniques used here may also help share other experimental codepoints, but that issue is out of scope for this document.

2. Conventions used in this document

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](#)].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying [RFC-2119](#) significance.

In this document, the characters ">>" preceding an indented line(s) indicates a compliance requirement statement using the key words listed above. This convention aids reviewers in quickly identifying or finding the explicit compliance requirements of this RFC.

3. TCP Experimental Option Structure

TCP options have the current common structure [[RFC793](#)], in which the first byte is the codepoint (Kind) and the second byte is the length of the option in bytes (Length):

```

0          1          2          3
01234567 89012345 67890123 45678901
+-----+-----+-----+-----+
| Kind  | Length |      ...      |
+-----+-----+-----+-----+
|      ...      |
+-----+
```

Figure 1 TCP Option Structure [[RFC793](#)]

This document extends the option structure for experimental codepoints (253, 254) with an experiment identifier (ExID), which is either 2 or 4 bytes in length. The ExID is used to differentiate different experiments, and is the first field after the Kind and Length, as follows:

Touch

Expires August 25, 2013

[Page 4]




Figure 2 TCP Experimental Option with a 16-bit ExID




Figure 3 TCP Experimental Option with a 32-bit ExID

>> Protocols requiring new TCP option codepoints that are not eligible for assigned values SHOULD use the existing TCP experimental option codepoints (253, 254) with ExIDs as described in this document.

>> All protocols using the TCP experimental option codepoints (253, 254) SHOULD use ExIDs as described in this document.

3.1. Selecting an ExID

ExIDs are selected at design time, when the protocol designer first implements or specifies the experimental option. ExIDs can be either 16-bits or 32-bits. In both cases, the value is stored in the header in network-standard (big-endian) byte order. ExIDs combine properties of IANA registered codepoints with "magic numbers".

ExIDs are registered with IANA using "first-come, first-served" priority based on the first two bytes. Those two bytes are thus sufficient to interpret which experimental option is contained in the option field.

The second two bytes serve as a "magic number". A magic number is a self-selected codepoint whose primary value is its unlikely collision with values selected by others. Magic numbers are used in other protocols, e.g., BOOTP [[RFC951](#)] and DHCP [[RFC2131](#)]. The magic number helps reduce the probability of a false positive collision with those who either do not register their experiment or who do not

implement this mechanism. Using the additional magic number bytes also helps the option contents have the same byte alignment in the TCP header as they would have if (or when) a conventional (non-experiment) TCP option codepoint is assigned.

3.2. Impact on TCP Option Processing

The ExID number is considered part of the TCP option, not the TCP option header. The presence of the ExID increases the effective option Length field by the size of the ExID. The presence of this ExID is thus transparent to implementations that do not support TCP options where it is used.

During TCP processing, ExIDs in experimental options are matched against the ExIDs for each implemented protocol. The remainder of the option is specified by the particular experimental protocol.

>> Experimental options that have ExIDs that do not match implemented protocols MUST be ignored.

The ExID mechanism must be coordinated during connection establishment, just as with any TCP option.

>> TCP ExID, if used in any TCP segment of a connection, MUST be present in TCP SYN segments of that connection.

>> TCP experimental option ExIDs, if used in any TCP segment of a connection, SHOULD be used in all TCP segments of that connection in which any experimental option is present.

Use of an ExID uses additional space in the TCP header and requires additional protocol processing by experimental protocols. Because these are experiments, neither consideration is a substantial impediment; a finalized protocol can avoid both issues with the assignment of a dedicated option codepoint later.

4. Reducing the Impact of False Positives

False positives occur where the ExID of one experiment matches the value of an option that does not use ExIDs or if two experiments select the same ExID. Such collisions can cause an option to be interpreted by the incorrect processing routine. Use of checksums or signatures may help an experiment use a shorter ExID while reducing the corresponding increased potential for false positives.

>> Experiments that are not robust to ExID false positives SHOULD implement other detection measures, such as checksums or minimal digital signatures over the experimental options they support.

5. Migration to Assigned Options

Some experiments may transition from experiment, and become eligible for an assigned TCP option codepoint. This document does not recommend a specific migration plan to transition from use of the experimental TCP options/ExIDs to use of an assigned codepoint.

However, once an assigned codepoint is allocated, use of an ExID represents unnecessary overhead. As a result:

>> Once a TCP option codepoint is assigned to a protocol, that protocol SHOULD NOT continue to use an ExID as part of that assigned codepoint.

This document does not recommend whether or how an implementation of an assigned codepoint can be backward-compatible with use of the experimental codepoint/ ExID.

However, some implementers may be tempted to include both the experimental and assigned codepoint in the same segment, e.g., in a SYN to support backward-compatibility during connection establishment. This is a poor use limited resources and so to ensure conservation of the TCP option space:

>> A TCP segment MUST NOT contain both an assigned TCP option codepoint and a TCP experimental option codepoint for the same protocol.

Instead, a TCP that intends backward compatibility might send multiple SYNs with alternates of the same option and discard all but the most desired successful connection. Although this approach may resolve more slowly or require additional effort at the endpoints, it is preferable to excessively consuming TCP option space.

6. Security Considerations

The mechanism described in this document is not intended to provide (nor does it weaken existing) security for TCP option processing.

7. IANA Considerations

This document calls for IANA to create a new TCP experimental option Experiment Identifier (ExID) registry.

That registry should allow 16-bit and 32-bit entries, where entries are "first-come, first-served" on the first two bytes of the value in network-standard byte order (big endian), in which the entry should indicate the entire ExID value. Known overlapping uses - whether of the first-come portion or the entire value - should also be listed and highlighted as collisions.

IANA should impose no requirements on making a registration other than indicating the desired codepoint and providing a point of contact. A short description or acronym for the use is desired, but should not be required.

8. References

8.1. Normative References

- [RFC793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), Sep. 1981.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4727] Fenner, B., "Experimental Values in IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers", [RFC 4727](#), Nov. 2006.

8.2. Informative References

- [Bi11] Bittau, A., D. Boneh, M. Hamburg, M. Handley, D. Mazieres, Q. Slack, "Cryptographic protection of TCP Streams (tcpcrypt)", work in progress, [draft-bittau-tcp-crypt-03](#), Sep. 3, 2012.
- [Ed11] Eddy, W., "Additional TCP Experimental-Use Options", work in progress, [draft-eddy-tcpm-addl-exp-options-00](#), Aug. 16, 2011.
- [IANA] IANA web pages, <http://www.iana.org/>
- [RFC951] Croft, B., J. Gilmore, "BOOTSTRAP PROTOCOL (BOOTP)", [RFC 951](#), Sept. 1985.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), Oct. 1996.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), Mar. 1997.

- [RFC2780] Bradner, S., V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", [BCP 37](#), [RFC 2780](#), Mar. 2000.
- [RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", [BCP 82](#), [RFC 3692](#), Jan. 2004.
- [RFC5226] Narten, T., H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC6013] Simpson, W., "TCP Cookie Transactions (TCPCT)", [RFC 6013](#), Jan. 2011.
- [Si11] Simpson, W., "TCP Cookie Transactions (TCPCT) Sockets Application Program Interface (API)", work in progress, [draft-simpson-tcpct-api-04](#), Apr. 7, 2011.

9. Acknowledgments

This document was motivated by discussions on the IETF TCPM mailing list and by Wes Eddy's proposal [[Ed11](#)]. Yoshifumi Nishida, Pasi Sarolathi, and Michael Scharf provided detailed feedback.

This document was prepared using 2-Word-v2.0.template.dot.

Authors' Addresses

Joe Touch
USC/ISI
4676 Admiralty Way
Marina del Rey, CA 90292-6695 U.S.A.

Phone: +1 (310) 448-9151
Email: touch@isi.edu