

Network Working Group
Internet-Draft
Updates: [4727](#) (if approved)
Intended status: Standards Track
Expires: February 17, 2012

W. Eddy
MTI Systems
August 16, 2011

Additional TCP Experimental-Use Options
draft-eddy-tcpm-addl-exp-options-00

Abstract

There have been multiple issues with the allocation of TCP option kind numbers recently. Two of these issues, which this document attempts to address, are that there were only a small number of options reserved by [RFC 4727](#) for experiment and test use in the [RFC 3692](#) style to begin with, and both of these have been used in shipping products. This impacts the ability of other research and experimental efforts to develop and test running code since registration of other option numbers requires either IESG Approval or Standards Action. This document proposes designation of additional experimental options in the IANA registry for TCP Option Kind Numbers, intended to resolve the possible barriers to using the existing [RFC 3962](#) experimental-use options.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may not be modified, and derivative works of it may not be created, except to format it for publication as an RFC and to translate it into languages other than English.

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 February 17, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the

Internet-Draft

TCP Experimental Options

August 2011

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.

Table of Contents

1.	Introduction	3
2.	Additional Experimental-Use Options	4
3.	Rationale	4
4.	Security Considerations	5
5.	IANA Considerations	5
6.	References	6
6.1.	Normative References	6
6.2.	Informative References	6
	Author's Address	6

1. Introduction

TCP options are a fundamental mechanism for extending and enhancing TCP's functionality. In the past, the addition of TCP options (e.g. Window Scale, Timestamp [[RFC1323](#)], and Selective Acknowledgements [[RFC2018](#)]) has supported the protocol's evolution and helped its applicability to expand as the Internet and types of links available grew.

However, there has been significant confusion with regards to how TCP option kind numbers are managed. This is, frankly, dangerous to the Internet, if it persists. There is a limited pool of options and due to misunderstandings the usable portion of this pool has shrunk to an unknown extent.

Registration of TCP option kind numbers is a function of the IANA [[RFC2780](#)]. Values are assigned following either (1) IESG Approval, or (2) Standards Action process, which are defined in [[RFC2434](#)]. Some vendors have not followed these procedures and simply shipped products using option kind numbers chosen themselves. This poisons the pool of options available, as it potentially causes conflicts if IANA later registers those same kind numbers for a use that followed the proper registration process. This has been recognized as a mistake, and vendors have expressed a desire to avoid it in the future and are working towards possible transition of such products to registered options numbers.

Two TCP option numbers have been designated for experimental use [[RFC4727](#)], which are not intended to be used in general deployments or enabled by default in products or other general releases unless explicitly enabled by an end-user [[RFC3692](#)].

Unfortunately, at least one vendor intending to avoid shipping its products using unregistered option numbers, actually shipped products using the experimental-use numbers. These numbers are being used by some deployed middleboxes and the impacts to other people trying to

use the same kind numbers for other purposes is not broadly understood, especially since the presence of such middleboxes on a path may be unknown a priori.

A recent TCP research effort testing running code over the Internet that would have been a perfect candidate for using the experimental-use numbers shied away from this due to the deployed middlebox issue and chose to improperly use yet more unregistered TCP option kind numbers.

Another recent issue is that with multiple ongoing efforts to extend TCP, there may be implementations that integrate a number of

extensions requiring experimental-use options. Two kind numbers may not be sufficient for such cases, and adding sub-kind identifiers within the option payload may be complex or even impossible.

This document attempts to mitigate the situation and remove excuses for such instances in the future by requesting IANA to register a greater number of TCP experimental-use options that would also follow the [RFC 3692](#) spirit for their intended use.

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.](#) Additional Experimental-Use Options

This document proposes to create an additional sixteen experimental-use TCP option kinds in the spirit of [RFC 3692](#). As this may seem like a large number compared to the current two that [RFC 4727](#) requested, some rationale is provided in the next section.

Of these sixteen option kinds, the option-length field for all of them will be defined as variable, but in all cases will hold a value of at least 2 in order to account for the kind and length fields.

The option kind numbers allocated should be contiguous in order to support potential ease of updating filter rules and other databases used in firewalls and other middleboxes, as well as various other software tools for packet analysis and other uses.

[3.](#) Rationale

There are only 8 bits that comprise a TCP Option Kind field, lending 256 possible unique codepoints. Of these, there are the 2 identified in [RFC 4727](#) for experimental-use and 19 with registrations that are currently not identified as obsolete (historic and currently unused) or unassigned due to release of prior registrations. Of these, several are not known to be in general use and could likely be reaped if needed. Additionally, 11 kind numbers have been identified as obsolete or unassigned due to registration being released, and 6 more are known to be deployed without proper IANA assignment. One further protocol under development in the IETF (Multipath TCP) requires an IANA option kind assignment yet-to-be-made.

This leaves 217 option kind values that both have never been registered and are not known to have been under deployment without registration. Even though this document proposes to claim 16 of

Eddy

Expires February 17, 2012

[Page 4]

Internet-Draft

TCP Experimental Options

August 2011

these values for experimental-use, there will still be 201 option kind values seemingly fully available, which represents over 78% of the option kind numbers. Based on TCP's existence for several decades without even using a quarter of the available options space, the remaining pool of kind numbers should be sufficient for many more decades to come.

Further, 16 option numbers for experimental use should be more than sufficient by a factor of 2 to 4 in order to permit implementing and testing combinations of experimental TCP extensions that do not yet have their own registered option kind numbers. This is especially true as recently Multipath TCP design has set an example for using a sub-kind / subtype field in order to avoid requiring multiple kind numbers from the TCP registry. This practice could be reused by future similar extensions making extensive use of TCP options.

[4.](#) Security Considerations

This document creates no additional security considerations for TCP implementations.

Firewalls and other network devices that aggressively filter unrecognized TCP options may cause difficulties in using the new experimental-use kind numbers defined by this document. Managers and vendors of such firewalls should reconsider whether such filtering is necessary or useful as this practice represents a major impediment to innovation in TCP.

[5.](#) IANA Considerations

This document requests that IANA allocate sixteen contiguous TCP option values for experimental-use in the spirit of [RFC 3692](#), which will be described in the registry as:

- o Length: N
- o Meaning: [RFC3692](#)-style Experiment - MUST NOT be used by default in shipping products, or other uncontrolled wide-scale deployments outside of an experimental context
- o Reference: (this document's RFC number, to be filled in)

Allocation from the rear of the available reserved space adjacent below the two existing experimental-use options (253 and 254) is desirable.

Eddy

Expires February 17, 2012

[Page 5]

Internet-Draft

TCP Experimental Options

August 2011

[6.](#) References

[6.1.](#) Normative References

- [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](#), November 2006.

[6.2.](#) Informative References

- [RFC1323] Jacobson, V., Braden, B., and D. Borman, "TCP Extensions for High Performance", [RFC 1323](#), May 1992.

- [RFC2018] Mathis, M., Mahdavi, J., Floyd, S., and A. Romanow, "TCP Selective Acknowledgment Options", [RFC 2018](#), October 1996.
- [RFC2434] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 2434](#), October 1998.
- [RFC2780] Bradner, S. and V. Paxson, "IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers", [BCP 37](#), [RFC 2780](#), March 2000.
- [RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", [BCP 82](#), [RFC 3692](#), January 2004.

Author's Address

Wesley M. Eddy
MTI Systems
MS 500-ASRC
NASA Glenn Research Center
21000 Brookpark Road
Cleveland, OH 44135
USA

Phone: +1-216-433-6682
Email: wes@mti-systems.com