Network Working Group Internet-Draft Intended status: Experimental Expires: June 11, 2017

TCP Capability Option draft-boucadair-tcpm-capability-option-01

Abstract

This document defines an experimental TCP option that can be used to negotiate a set of options that are supported by two TCP endpoints. The main motivation for designing this option is the optimization of the option space for SYN segments.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of <u>BCP 78</u> and <u>BCP 79</u>.

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 <u>http://datatracker.ietf.org/drafts/current/</u>.

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 11, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents (<u>http://trustee.ietf.org/license-info</u>) 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

$\underline{1}$. Introduction	•	•	•		•	•	•	•	•	•	•			•	•		<u>2</u>
<u>1.1</u> . Applicability																	<u>3</u>
<u>1.2</u> . Requirements Language																	<u>3</u>
$\underline{2}$. TCP Capability Option	•			•							•			•		•	<u>3</u>
$\underline{3}$. Security Considerations .																	<u>4</u>
4. IANA Considerations																	<u>4</u>
<u>5</u> . References																	<u>4</u>
<u>5.1</u> . Normative References																	<u>4</u>
<u>5.2</u> . Informative Reference	S																<u>5</u>
Authors' Addresses																	<u>5</u>

1. Introduction

TCP ([RFC0793]) can be extended by defining new options. Because of the presence of NATs, firewalls, and other types of flow-aware service functions between two TCP endpoints, means to ensure that both endpoints support a given option, and that the paths used to forward traffic between them do not involve nodes that strip or alter the content of the options. This is actually achieved by negotiating the support of a given option during the 3WHS stage.

Typically, an option is included in the SYN message to inform the remote TCP peer that the originating TCP peer is "X"-capable (that is, it supports feature "X"). Upon receipt of the SYN message, and if no intermediary node has stripped that option, the remote peer will echo that option in a SYN/ACK if and only if it is also "X"-capable. Feature "X" can then be used if the SYN/ACK message that is received by the originating peer still carries the "X"-capable; the option must then be included in the ACK.

For example, several TCP extensions have been designed with this design rationale in mind, e.g., SYN-EOS [<u>I-D.touch-tcpm-tcp-syn-ext-opt</u>], EDO [<u>I-D.ietf-tcpm-tcp-edo</u>], SACK Permitted [<u>RFC2018</u>], MP_CAPABLE [<u>RFC6824</u>], etc. In the meantime, and given the limited TCP option space, it becomes more challenging to include all options in a single SYN message.

Several solutions have been proposed to solve that problem (e.g., [<u>I-D.touch-tcpm-tcp-syn-ext-opt</u>]) by means of generating a companion TCP message. This document proposes a solution that aims to optimize the required option space to facilitate the insertion of several "X"capable options.

[Page 2]

<u>1.1</u>. Applicability

The option is primarily designed for network configurations where the path between the TCP endpoints is managed (e.g., an MPTCP endpoint embedded in the CPE and the remote MPTCP endpoint is located in the network side; the paths between these endpoints are managed by the same administrative entity such as in [I-D.boucadair-mptcp-natfwfree-profile]).

A flow-aware device that removes the option will disable all the options that were included in the TCP Capability option. This is not supposed in the target deployment context for this option.

Some middleboxes may allow the TCP Capability option to pass through, but the individual options may be blocked. This is not supposed in the target deployment context for this option as those flow-aware functions are managed.

<u>1.2</u>. Requirements Language

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 <u>RFC 2119</u> [<u>RFC2119</u>].

2. TCP Capability Option

The format of the TCP Capability option is shown in Figure 1. The option follows the format defined in [<u>RFC0793</u>].

+----+ | Kind | Length | Kind 1| Kind 2 | +----+ : |Kind n-1| Kind n | +----+

Figure 1

"Kind" is to be assigned by IANA (see <u>Section 4</u>).

In order to conduct experiments with this option, a format that adheres to [RFC6994] is used (see Figure 2). EXID MUST be set to 0x0CA0 (3232).

+ •		+ •		-+	+	-+
	Kind	I	Length		ExID	
+ -		· + ·		-+	+	-+
:	Kind 1	Ι	Kind 2	I		:
+ -				-+		-+

Figure 2

"Kindi" carries the code point of an option that a TCP endpoint supports and which it would like to negotiate with the remote peer. One or multiple "Kindi" fields may be included.

An endpoint that wishes to enable the capabilities associated with one or multiple TCP options must include the corresponding "Kind" codes in a TCP Capability option, which is included in the initial SYN. If the remote peer supports at least one of the options carried in the TCP Capability option, it replies with a SYN/ACK that includes the TCP Capability option and which only carries the code points of the options it supports. These values are then echoed in the ACK message the originating peer sends back to the remote peer.

When replying to a SYN message that includes a TCP Capability option, the remote peer should preserve the same order of the "Kindi" fields (of course, those that are not supported won't be included).

<u>3</u>. Security Considerations

The security considerations are to be completed.

<u>4</u>. IANA Considerations

IANA is requested to record the following identifier in the "TCP Experimental Option Experiment Identifiers (TCP ExIDs)" registry:

Value	Description	Reference
0x0CA0	TCP Capability Option	[This document]

5. References

5.1. Normative References

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, <u>RFC 793</u>, DOI 10.17487/RFC0793, September 1981, <http://www.rfc-editor.org/info/rfc793>.

[Page 4]

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <u>BCP 14</u>, <u>RFC 2119</u>, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.
- [RFC6994] Touch, J., "Shared Use of Experimental TCP Options", <u>RFC 6994</u>, DOI 10.17487/RFC6994, August 2013, <<u>http://www.rfc-editor.org/info/rfc6994</u>>.

5.2. Informative References

[I-D.boucadair-mptcp-natfwfree-profile]

Boucadair, M., Jacquenet, C., Seite, P., Bonaventure, O., and D. Lingli, "An MPTCP Profile for NAT- and Firewall-Free Networks: Network-Assisted MPTCP Deployments", <u>draft-</u> <u>boucadair-mptcp-natfwfree-profile-00</u> (work in progress), July 2015.

[I-D.ietf-tcpm-tcp-edo]

Touch, J. and W. Eddy, "TCP Extended Data Offset Option", <u>draft-ietf-tcpm-tcp-edo-06</u> (work in progress), June 2016.

[I-D.touch-tcpm-tcp-syn-ext-opt]

Touch, J. and T. Faber, "TCP SYN Extended Option Space Using an Out-of-Band Segment", <u>draft-touch-tcpm-tcp-syn-</u> <u>ext-opt-05</u> (work in progress), October 2016.

- [RFC2018] Mathis, M., Mahdavi, J., Floyd, S., and A. Romanow, "TCP Selective Acknowledgment Options", <u>RFC 2018</u>, DOI 10.17487/RFC2018, October 1996, <<u>http://www.rfc-editor.org/info/rfc2018</u>>.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", <u>RFC 6824</u>, DOI 10.17487/RFC6824, January 2013, <<u>http://www.rfc-editor.org/info/rfc6824</u>>.

Authors' Addresses

Mohamed Boucadair Orange Rennes 35000 France

Email: mohamed.boucadair@orange.com

[Page 5]

Christian Jacquenet Orange Rennes 35000 France

Email: christian.jacquenet@orange.com