

MPTCP Working Group
Internet-Draft
Updates: [6824](#) (if approved)
Intended status: Experimental
Expires: December 18, 2015

O. Bonaventure
B. Hesmans
UCLouvain
M. Boucadair
Orange
June 16, 2015

Experimental Multipath TCP option
draft-bonaventure-mptcp-exp-option-00

Abstract

This document proposes a generic format for an experimental Multipath TCP option. This experimental option is expected to replace the option type that was reserved for private use in [[RFC6824](#)].

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 December 18, 2015.

Copyright Notice

Copyright (c) 2015 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.

Table of Contents

1.	Introduction	2
2.	The Experimental Multipath TCP option	3
3.	Security considerations	6
4.	IANA considerations	6
5.	Conclusion	6
6.	Acknowledgements	6
7.	References	7
7.1.	Normative References	7
7.2.	Informative References	7
	Authors' Addresses	8

[1.](#) Introduction

Multipath TCP is an extension to TCP [[RFC0793](#)] that was specified in [[RFC6824](#)]. Like other TCP extensions, Multipath TCP relies on TCP options to exchange information. To mitigate some forms of middlebox interference, Multipath TCP uses only one TCP option identified by Kind=30 [[RFC6824](#)]. TCP Option Kind=30 has a variable length and includes a subtype field that can support up to 16 different Multipath TCP options. As of this writing, 8 different Multipath TCP options are defined [[RFC6824](#)] :

- o subtype=0x0 : MP_CAPABLE
- o subtype=0x1 : MP_JOIN
- o subtype=0x2 : DSS
- o subtype=0x3 : ADD_ADDR
- o subtype=0x4 : REMOVE_ADDR
- o subtype=0x5 : MP_PRI0
- o subtype=0x6 : MP_FAIL
- o subtype=0x7 : MP_FASTCLOSE

In addition, subtype=0x8 is already reserved for a new version of the ADD_ADDR option [[I-D.ietf-mptcp-rfc6824bis](#)] and subtype=0xf is reserved for use in private testbeds [[RFC6824](#)]. There are thus less than six different Multipath TCP options that are available to extend the protocol.

The Multipath TCP options defined above cover the basic functionalities for Multipath TCP but previous experience with TCP

indicates that it can be expected that Multipath TCP implementors will propose new extensions that require a new option subtype. Some authors have already proposed such extensions

[[I-D.boucadair-mptcp-symmetric](#)] [[I-D.paasch-mptcp-syncookies](#)] [[I-D.xu-mptcp-prmp](#)]. Given the small number of available Multipath TCP options, it will be difficult to officially reserve them for all proposed utilisations even if experiments with these new Multipath TCP options would probably be useful.

This document proposes a shared use of the Multipath TCP option subtype=0xf following a similar approach as described in [[RFC6994](#)] for TCP. This approach is efficient as it encourages conducting new experiments while not consuming the limited option codepoint space.

2. The Experimental Multipath TCP option

The extensibility of TCP depends largely on the definition of new TCP options. As of this writing, 30 different TCP options have been assigned by the IANA among the 256 possible values. Two of these option types (253 and 254) have been reserved as experimental options in [[RFC4727](#)]. However, this document did not specify how these experimental options should be used. [[RFC6994](#)] clarified the situation by proposing two encodings for these experimental options to support a shared use of the same option kind by different experimental options. Basically, [[RFC6994](#)] proposes to use a 16 bits experiment identifier in the first two bytes of the option as shown in the figure below.

```

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


[[RFC6994](#)] extends this format to support 32 bits identifiers where the second 16 bits word is used to encode a magic number.

```

0          1          2          3
01234567 89012345 67890123 45678901
+-----+-----+-----+-----+
| Kind  | Length |      ExID      |
+-----+-----+-----+-----+
|  ExID (con't) | option contents...
+-----+-----+-----+-----+
```


The experimental TCP options defined in [[RFC6994](#)] could in theory be used over any TCP connection, including the subflows used by Multipath TCP. However, using them as such to support experimental features on a Multipath TCP connection would cause two practical problems. The first problem is that the utilisation of such an option should be negotiated during the three-way handshake by adding the experimental option in the SYN and SYN+ACK segments. Given the limited space for the options in these segments when Multipath TCP is used and other options are already used, this is not a valid approach. Second, middleboxes could accept the Multipath TCP option and drop or modify an experimental TCP option that is used to support an experimental Multipath TCP feature. This would increase the complexity of the middlebox handling algorithms inside a Multipath TCP implementation.

For the above two reasons, we propose the utilisation of an Experimental Multipath TCP option that is encoded as shown in the figure below.



Flags is a set of 4 flags:

```

+--+--+--+
|S|U|r|r|
+--+--+--+

```

where "rr" are for future assignment as additional flag bits. r bits MUST each be sent as zero and MUST be ignored on receipt.

The Experimental Multipath TCP option has a variable length. It is always of kind=30 and Subtype=0xf. The Experiment identifier is a 16 bits integer that shall be assigned by using the same procedure as defined in [[RFC6994](#)]. [[RFC6994](#)] also supports the utilisation of 32 bits Experiment Identifiers. These 32 bits identifiers are composed of a 16 bits identifier assigned on a FCFS basis and a 16 bits magic number. At this stage, we propose that the Experimental Multipath TCP option uses only the 16 bits Experiment Identifier.

The two high order flags that are included in the Experimental Multipath TCP option have the following semantics :

- o "S" flag (highest order bit) : This is the synchronising bit. When set to 1, it indicates that the host sending this option expects a reply from the remote host with an option having the same experiment identifier, but possibly containing other data.
- o "U" flag (second highest order bit) : When set to 1, this flag indicates that the experimental option was received by the sending host but it was unable to parse it.

The two low order flags are currently reserved for further use. They MUST be set to zero when sending and ignored upon reception.

To use the Experimental Multipath TCP option with a given experiment identifier over a Multipath TCP connection, the sending host must first verify the ability of the remote host to support this particular Experimental option. For this, it first sends in any valid TCP segment, including a duplicate acknowledgement, an Experimental Multipath TCP option with the "S" flag set. Upon reception of this option, the receiving host will verify whether it supports it. If yes, it shall return a TCP segment that contains the experimental option with the same identifier and the "S" and the "U" flags reset. This option may contain additional data depending on its semantics. If the receiving host does not recognise the experimental option that it has received, it shall return a TCP segment that contains the received experimental option with the "S" flag set to 0 and the "U" flag set to 1. Upon reception of this option, the sending host knows that the experimental option cannot be used over this Multipath TCP connection.

Once the utilisation of a specific Experimental Multipath TCP option has been negotiated over a Multipath TCP connection, both hosts can use this experimental Multipath TCP option. If a host receives an Experimental Multipath TCP option with the "U" flag set to 0 that it does not support or contains information that it cannot parse, it shall return the exact option that it received with the "U" flag set to indicate the error to the remote host. If an invalid option is received with the "U" flag set to 0, it must be silently discarded.

Future documents specifying new experimental MPTCP options should specify the extract semantic of the Subtype-specific data and whether additional validation operations are to be followed at both sides. Note, data can be included in an experimental data concurrently with the capability check (S/U).

As a side effect of the use of a shared codepoint, an overhead of the 2 octets will be consumed in the option space for each experimental MPTCP option. This overhead should be taken into account when drawing conclusions related to the co-existence of multiple MPTCP

experimental options or, even worse, between experimental MPTCP options and experimental TCP options.

3. Security considerations

This document does not modify the security considerations for Multipath TCP.

4. IANA considerations

This document proposes to modify the Multipath TCP option subtype 0xf that was reserved for utilisation in private testbed to support the Experimental Multipath TCP option.

The same considerations as for [[RFC6994](#)] apply. IANA should create a "Multipath TCP Experimental Option Identifiers (MPTCP ExIDs)" registry. This registry contains the 16 bits ExIDs and a reference (description, document pointer, or assignee name and e-mail contact) for each entry. MPTCP ExIDs are assigned on a First Come, First Served (FCFS) [[RFC5226](#)].

IANA will advise applicants of duplicate entries to select an alternate value, as per typical FCFS processing.

IANA will record known duplicate uses to assist the community in both debugging assigned uses as well as correcting unauthorized duplicate uses.

IANA should impose no requirement 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.

5. Conclusion

This document has proposed a new experimental Multipath TCP option that is intended to replace the already defined, but apparently currently unused, Multipath TCP option reserved for private testbed. This new Experimental Multipath TCP option will enable implementors to perform experiments with Multipath TCP options without using the limited Multipath TCP option subtypes.

6. Acknowledgements

This work was partially supported by the FP7-Trilogy2 project. This document was heavily inspired by [[RFC6994](#)].

[7.](#) References

[7.1.](#) Normative References

- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", [RFC 6824](#), January 2013.

[7.2.](#) Informative References

- [I-D.boucadair-mptcp-symmetric]
Boucadair, M. and C. Jacquenet, "An Extension to MPTCP for Symmetrical Sub-Flow Management", [draft-boucadair-mptcp-symmetric-02](#) (work in progress), March 2015.
- [I-D.ietf-mptcp-rfc6824bis]
Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", [draft-ietf-mptcp-rfc6824bis-04](#) (work in progress), March 2015.
- [I-D.paasch-mptcp-syncookies]
Paasch, C., Biswas, A., and D. Haas, "Making Multipath TCP robust for stateless web servers", [draft-paasch-mptcp-syncookies-00](#) (work in progress), April 2015.
- [I-D.xu-mptcp-prmp]
Xu, C., Huang, H., Zhang, H., Xiong, C., and L. Zhu, "Multipath Transmission Control Protocol (MPTCP) Partial Reliability Extension", [draft-xu-mptcp-prmp-00](#) (work in progress), April 2015.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [RFC4727] Fenner, B., "Experimental Values In IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers", [RFC 4727](#), November 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC6994] Touch, J., "Shared Use of Experimental TCP Options", [RFC 6994](#), August 2013.

Authors' Addresses

Olivier Bonaventure
UCLouvain

Email: Olivier.Bonaventure@uclouvain.be

Benjamin Hesmans
UCLouvain

Email: Benjamin.Hesmans@uclouvain.be

Mohamed Boucadair
Orange

Email: mohamed.boucadair@orange.com

