

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: May 28, 2017

M. Boucadair
C. Jacquenet
Orange
November 24, 2016

Negotiating the Maximum Number of Multipath TCP (MPTCP) Subflows
draft-boucadair-mptcp-max-subflow-03

Abstract

This document specifies an experimental Multipath TCP (MPTCP) option that is meant to negotiate the maximum number of subflows that can be established and maintained for a given MPTCP connection. The purpose is to minimize any possible performance degradation that can be induced by a possibly large number of establishment requests for additional subflows if the remote endpoint is not appropriately dimensioned to handle such requests.

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 [RFC 2119](#) [[RFC2119](#)].

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 May 28, 2017.

Copyright Notice

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

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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
1.1.	Experiment Goals	3
2.	Maximum Subflows MPTCP Option (MP_MAXSUBFLOW)	3
3.	Behavior	3
4.	Security Considerations	4
5.	IANA Considerations	4
6.	References	4
6.1.	Normative References	4
6.2.	Informative References	5
	Authors' Addresses	5

[1.](#) Introduction

This document specifies a Multipath TCP (MPTCP) [[RFC6824](#)] option that is meant to indicate to a remote peer the maximum number of subflows that can be established within a single MPTCP connection. If the remote peer honors the indication provided in this option, any performance degradation induced by a possibly abusive setup of additional subflows that exceed the said maximum becomes unlikely.

This option mainly targets MPTCP deployments within a single administrative domain such as those MPTCP designs meant to achieve load-balancing, for example. The use of this option contributes to the harmonization of node configuration within an administrative domain, so that an optimal number of subflows is maintained by involved nodes independently of their actual performance capabilities. This option can be used in other deployment scenarios. It is out of scope of this document to identify what are such scenarios.

Section 3.4 of [[I-D.ietf-mptcp-experience](#)] includes a discussion on

the MPTCP session manager issues. That document states that creating subflows between multihomed clients and servers may sometimes lead to operational issues. The present document targets to soften these issues.

1.1. Experiment Goals

Experiments based upon the MPTCP option described in this document are meant to help operators optimize their MPTCP design and operational procedures, by tweaking some MPTCP parameters such as the number of subflows to be associated with a given MPTCP connection. Experimenting with this MPTCP option should also help assess whether this option can be used to propagate MPTCP-related optimization parameters (derived from the number of concurrent subflows associated to each MPTCP connection) that can be configured in a node that is responsible for aggregating MPTCP connections established with upstream nodes.

2. Maximum Subflows MPTCP Option (MP_MAXSUBFLOW)

The MP_MAXSUBFLOW option (Figure 1) follows the shared experimental format defined in [\[I-D.ietf-mptcp-rfc6824bis\]](#).

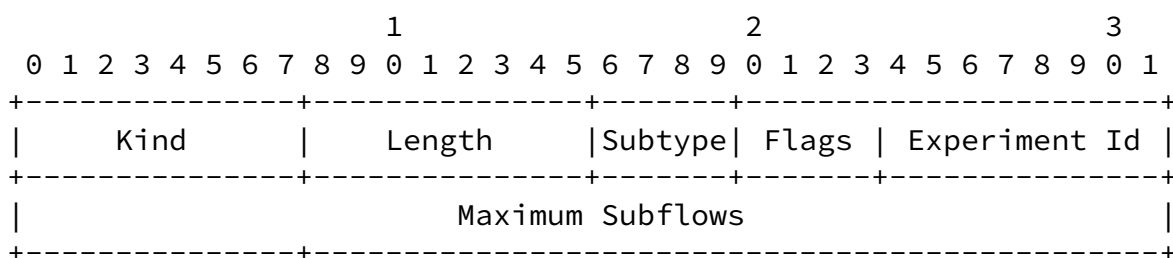


Figure 1: Option Format

The meaning of "Kind", "Length", "Subtype", and "Flags" (especially 'S' and 'U' flags) are exactly the same as defined in Section 3.7 of [\[I-D.ietf-mptcp-rfc6824bis\]](#).

Experiment ID MUST be set to 0xAA (see [Section 5](#)).

The "Maximum Subflows" field indicates the maximum number of concurrent subflows that can be maintained by a given MPTCP endpoint

for each MPTCP connection established by or with this endpoint. The value of this field MUST be strictly greater than zero.

3. Behavior

The MP_MAXSUBFLOW option may be used by a TCP endpoint to indicate to its corresponding peer the maximum number of subflows that it can maintain per MPTCP connection.

How an MPTCP endpoint determines the maximum number of concurrent subflows it supports (i.e., the value it indicates in an MP_MAXSUBFLOW option) is implementation-specific. For example, an

implementation may support a configuration parameter and/or use a default value.

If two peers (T1 and T2) indicate the maximum number of concurrent subflows per connection they can maintain, then they MUST NOT maintain more than "MIN(MAX_SUBFLOW(T1), MAX_SUBFLOW(T2))" concurrent subflows.

The absence of the MP_MAXSUBFLOW option in an MPTCP control message issued by a MPTCP endpoint is an indication that this endpoint can instantiate any number of subflows per MPTCP connection.

4. Security Considerations

MPTCP-related security considerations are documented in [[RFC6824](#)] and [[RFC7430](#)].

Establishing a large number of subflows may exhaust the resources of an MPTCP implementation (especially, MPTCP servers). MPTCP implementations SHOULD rate-limit the number of concurrent subflows per MPTCP connection. This policy is meant to prevent DoS attacks that would result in creating additional subflows to exhaust the remote peer resources (typically, an MPTCP server).

5. IANA Considerations

This document uses the experiment Id (0xAA). This Id is not registered yet given that no registry is maintained by IANA for this purpose.

[6.](#) References

[6.1.](#) Normative References

- [I-D.ietf-mptcp-rfc6824bis]
Ford, A., Raiciu, C., Handley, M., Bonaventure, O., and C. Paasch, "TCP Extensions for Multipath Operation with Multiple Addresses", [draft-ietf-mptcp-rfc6824bis-07](#) (work in progress), October 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", [RFC 6824](#), DOI 10.17487/RFC6824, January 2013, <<http://www.rfc-editor.org/info/rfc6824>>.

[6.2.](#) Informative References

- [I-D.ietf-mptcp-experience]
Bonaventure, O., Paasch, C., and G. Detal, "Use Cases and Operational Experience with Multipath TCP", [draft-ietf-mptcp-experience-07](#) (work in progress), October 2016.
- [RFC7430] Bagnulo, M., Paasch, C., Gont, F., Bonaventure, O., and C. Raiciu, "Analysis of Residual Threats and Possible Fixes for Multipath TCP (MPTCP)", [RFC 7430](#), DOI 10.17487/RFC7430, July 2015, <<http://www.rfc-editor.org/info/rfc7430>>.

Authors' Addresses

Mohamed Boucadair
Orange

Rennes 35000
France

Email: mohamed.boucadair@orange.com

Christian Jacquenet
Orange
Rennes 35000
France

Email: christian.jacquenet@orange.com