

MPTCP Working Group
Internet-Draft
Intended status: Informational
Expires: January 9, 2020

V. Tran
O. Bonaventure
Universite catholique de Louvain
July 08, 2019

Multipath TCP Subflow Rate Limit Option
draft-hoang-mptcp-sub-rate-limit-00

Abstract

This document defines a new MPTCP Option that enables hosts to request their peers to limit the maximum transfer rate on a per-subflow basis.

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 <https://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 January 9, 2020.

Copyright Notice

Copyright (c) 2019 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 (<https://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. Terminology	3
3. The Multipath TCP Subflow Rate Limit (SRL) Option	3
3.1. Option Format	3
3.2. SRL Option and Local Policies	4
4. Implementation and Interoperability	4
5. Security Considerations	4
6. IANA Considerations	5
7. Acknowledgements	5
8. References	5
8.1. Normative References	5
8.2. Informative References	5
Authors' Addresses	6

1. Introduction

Multipath TCP [RFC6824] [I-D.ietf-mptcp-rfc6824bis] is used in various use cases [RFC8040]. In several situations, a Multipath TCP host would like its peer to limit the sending rate over a specific subflow.

It is common for mobile clients to have limited cellular data subscription. Even if this is not the case, mobile network operators may still silently throttle the networking capacity of the customers who have used up a large amount of cellular data. A good LTE or 5G connection running at full speed during less than a few hours could consume the entire monthly budget cellular quota of many users. This is even more important when the mobile clients are roaming abroad where the monetary cost for cellular data can be very high. A common scenario is that mobile users want to limit the monetary cost of using cellular networks or to avoid running out of their mobile data quota. Smartphones can easily rate limit their upstream bandwidth, but unfortunately, most smartphone applications mainly receive data. For these applications, a rate limit must happen on the server side. This rate limit could be enforced by the application, e.g. by selecting a specific video coding scheme, but applying it at the transport layer would be more generic and could be done from the system level automatically.

As discussed on the multipathtcp IETF mailing list [paasch_mptcpwg_2019], this rate-control mechanism can also be used when a client wants to inform a server to close a subflow gracefully by requesting a zero transfer rate. Though the client may send a TCP-RST on this subflow instead, in-flight data would be lost and must be reinjected over other subflows. Another solution is to send an MP_PRIO option to the sender to put the cellular subflow into

backup mode, but this request could be overridden by the sender's local policy.

2. Terminology

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in BCP 14, RFC 2119 [RFC2119].

3. The Multipath TCP Subflow Rate Limit (SRL) Option

This document proposes a Subflow Rate Limit option that indicates a maximum receive rate for the subflow it is sent. Like other MPTCP options, this option is not sent reliably. Hosts SHOULD resend is several times, but not more frequently than once per second.

3.1. Option Format

The format of the SRL option is depicted in Fig. Figure 1:




Figure 1: MPTCP Subflow Rate Limit Option Format

All reserved bits MUST be set to zero.

In the SRL option, the Requested Rate (32 bits) is specified in IEEE 754 binary32 format (single-precision binary floating-point format). The same format is used in [RFC3630] to specify the maximum bandwidth of a link. This format is as follows:




Figure 2: IEEE binary32 format

- o Sign bit (1 bit) MUST be set to zero.
- o Exponent (8 bits) is the exponent base 2 in "excess 127" notation.
- o Fraction (23 bits) is the mantissa - 1, with an implied binary point in front of it.

Thus, the above represents the value:

$$(-1)**(S) * 2**(Exponent-127) * (1 + Fraction)$$

The unit of this value is Kilobits per second (Kbps).

3.2. SRL Option and Local Policies

Note that the SRL option is an indicative value. Upon reception of this option, the receiver SHOULD set the maximum rate on the subflow over which the option was received.

Like all Multipath TCP options, the SRL Option is exchanged without any protection from TCP's reliability mechanisms. Therefore, implementations MUST NOT assume that it is transferred reliably. Implementations that use the SRL option can transmit the SRL option at any time. Since the utilisation of this option is not negotiated during the connection handshake, a host MUST NOT send more than three SRL options on a connection where it has not received any SRL option.

4. Implementation and Interoperability

Implementations MAY use various mechanisms to implement the rate control policy, for example using TCP Pacing or clamping the subflow congestion window.

5. Security Considerations

Since the SRL option is neither encrypted nor authenticated, on-path attackers and middleboxes could remove, add or modify the SRL option on observed Multipath TCP connections. However, manipulating this option does not open new attacks compared to the ones documented in [RFC6181] [RFC7430].

For example, an on-path middle man could insert an option to throttle the rate on a subflow to nearly zero, effectively stalling the subflow. However, if an attacker has that capability, it could instead drop all packets or inject the TCP-RST/MP-FASTCLOSE.

On the other hand, on-path middleboxes may increase the rate-limit value in the exchanged option to a very high value. This effectively has the same effect as filtering out the option.

6. IANA Considerations

IANA is requested to assign an MPTCP option subtype for the SRL option from the "MPTCP Option Subtypes" available at <https://www.iana.org/assignments/tcp-parameters/tcp-parameters.xhtml>

7. Acknowledgements

This work was supported by the ARC-SDN project and the Walinnov MQUIC project, No. 1810018.

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6181] Bagnulo, M., "Threat Analysis for TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6181, DOI 10.17487/RFC6181, March 2011, <<https://www.rfc-editor.org/info/rfc6181>>.
- [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, <<https://www.rfc-editor.org/info/rfc6824>>.
- [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, <<https://www.rfc-editor.org/info/rfc7430>>.

8.2. Informative 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-18 (work in progress), June 2019.

[paasch_mptcpwg_2019]

Paasch, C., "Regarding rate control at a subflow level",
May 2019,
<[https://mailarchive.ietf.org/arch/msg/multipathtcp/
fyhIpWnXCr0sImRECGh8LxlzWdc](https://mailarchive.ietf.org/arch/msg/multipathtcp/fyhIpWnXCr0sImRECGh8LxlzWdc)>.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering
(TE) Extensions to OSPF Version 2", RFC 3630,
DOI 10.17487/RFC3630, September 2003,
<<https://www.rfc-editor.org/info/rfc3630>>.

[RFC8040] Bierman, A., Bjorklund, M., and K. Watsen, "RESTCONF
Protocol", RFC 8040, DOI 10.17487/RFC8040, January 2017,
<<https://www.rfc-editor.org/info/rfc8040>>.

Authors' Addresses

Viet-Hoang Tran
Universite catholique de Louvain

Email: hoang.tran@uclouvain.be

Olivier Bonaventure
Universite catholique de Louvain
Pl. Ste Barbe, 2
Louvain-la-Neuve 1348
Belgium

Email: olivier.bonaventure@uclouvain.be