

MPTCP Working Group
Internet-Draft
Expires: April 21, 2016

F. Duchene
UCLouvain
C. Paasch
Apple
O. Bonaventure
UCLouvain
October 19, 2015

Multipath TCP MIB
draft-duchene-mptcp-mib-00

Abstract

This memo proposes a simple Management Information Base (MIB) that gathers statistics and counters about the operation of a Multipath TCP implementation.

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 April 21, 2016.

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 MPTCP Group 2
- 3. IANA considerations 8
- 4. Conclusion 9
- 5. Acknowledgements 9
- 6. Normative References 9
- Authors' Addresses 9

1. Introduction

Multipath TCP is an extension to TCP [RFC0793] that was specified in [RFC6824]. Multipath TCP allows hosts to use multiple paths to send and receive the data belonging to one connection. For this, a Multipath TCP is composed of several TCP connections that are called subflows.

This memo proposes a simple Management Information Base (MIB) that gathers statistics and counters about the operation of a Multipath TCP implementation. They are designed to give a better understanding of Multipath TCP operations. In particular, the MIB covers the different failure conditions that might occur and would trigger a fallback to regular TCP of the MPTCP connection or the failure of a subflow. These failures have important operational implications. Further, several counters are defined to track the transmission and reception of data at the MPTCP-layer. These counters might help to understand the performance of MPTCP.

2. The MPTCP Group

mptcp OBJECT IDENTIFIER ::= { mib-2 TBD } - TBD

mptcpPassiveConn OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP connections that made a direct
    transition to the ESTABLISHED state from the
    SYN-RECEIVED + MP_CAPABLE state. In this MIB,
    an MPTCP connection
    is defined as a TCP connection where both the
    SYN and the SYN + ACK segments
    contained the MP_CAPABLE options."
 ::= { mptcp 1 }
```

mptcpActiveConn OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP connections that made a direct
    transition to the ESTABLISHED state from the
    SYN-SENT + MP_CAPABLE state."
 ::= { mptcp 2 }
```

mptcpPassiveCsumEnabled OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP connections that made a direct
    transition to the ESTABLISHED state from the
    SYN-RECEIVED + MP_CAPABLE state with the DSS CHECKSUM
    enabled."
 ::= { mptcp 3 }
```

mptcpActiveCsumEnabled OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP connections that made a direct
    transition to the ESTABLISHED state from the
    SYN-SENT + MP_CAPABLE state with the DSS CHECKSUM
    enabled."
 ::= { mptcp 4 }
```

mptcpPassiveRemovedSubflows OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP subflows that made a transition to the
    CLOSED state from the ESTABLISHED state upon reception of
    a segment containing REM_ADDR or RST."
 ::= { mptcp 5 }
```

mptcpActiveRemovedSubflows OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP subflows that made a transition to the
    CLOSED state from the ESTABLISHED state upon emission of
    a segment containing REM_ADDR or RST."
 ::= { mptcp 6 }
```

mptcpPassiveAddedSubflows OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP subflows that made a direct
    transition to the ESTABLISHED state from the
    SYN-RECEIVED + MP_JOIN state."
 ::= { mptcp 7 }
```

mptcpActiveAddedSubflows OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP subflows that made a
     transition to the ESTABLISHED state from the
     SYN-SENT + MP_JOIN state through the
     PRE_ESTABLISHED state."
 ::= { mptcp 8 }
```

mptcpFailedToEstablishInitialSubflows OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of initial MPTCP subflows (i.e. the initial SYN
     segment contained the MP_CAPABLE option) that could not
     transition to the ESTABLISHED state from the SYN-RECEIVED
     or SYN-SENT states. The reason being one of:
     - the SYN+ACK didn't contain a MP_CAPABLE
     - the first ACK didn't contain a DATA_ACK or the first
     data-segment did not contain a DSS mapping
     - 4-way handshake didn't complete (SYN+ACK or ACK not received)
     Given these reasons, a connection could not get established or fell
     back to regular TCP. They are most likely due to middleboxes
     interfering with the connection."
 ::= { mptcp 9 }
```

mptcpFailedToEstablishSubflows OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of non-initial MPTCP subflows (i.e. subflows
    started with a SYN containing the MP_JOIN option)
    that could not transition to the ESTABLISHED state
    from the SYN-RECEIVED or SYN-SENT states. The reason being:
    - the SYN+ACK didn't contain a MP_JOIN
    - the first ACK didn't contain a DATA_ACK or the first
    data-segment did not contain a DSS mapping
    - The connection had already failed back to TCP
    - the 4-way handshake didn't complete (SYN+ACK or ACK
    not received)
    Given, these reasons a subflow could not get established.
    They are most likely due to middleboxes interference."
 ::= { mptcp 10 }
```

mptcpFallbackEstablishedConnections OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP connections that fell back to regular TCP
    while being already ESTABLISHED. The reason being one of:
    - Reception of more than a window worth of data without DSS
    - Reception of a segment with an incorrect DSS checksum
    This happens when a middlebox is interfering with the data
    flow after the connection has been successfully established."
 ::= { mptcp 11 }
```

mptcpOtherFailures OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP connections were an invalid segment was
    received
    - Bad DSS-mapping (aka, the specified DSS-mapping does not map the
    TCP sequence numbers)
    - All other possible failures"
 ::= { mptcp 12 }
```

mptcpInvalidJoinReceived OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "segments"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of SYN+MP_JOIN segments that were received but
    discarded due to :
        - Error in the HMAC
        - Token not found"
 ::= { mptcp 13 }
```

mptcpFailRX OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "connections"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MPTCP connections where the remote host
    initiated a fallback. This counter is triggered by the
    reception of a MP_FAIL."
 ::= { mptcp 14 }
```

mptcpRetrans OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "segments"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of segments that where retransmitted
    at the MPTCP (meta) level, including all the types
    of reinjections."
 ::= { mptcp 15 }
```

mptcpFastCloseRX OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "segments"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MP_FASTCLOSE received."
 ::= { mptcp 16 }
```

mptcpFastCloseTX OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "segments"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of MP_FASTCLOSE emitted."
 ::= { mptcp 17 }
```

mptcpReceivedInOrder OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "segments"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of segments that were received in order at the
    MPTCP (meta) level."
 ::= { mptcp 18 }
```

mptcpReceivedOutOfOrder OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "segments"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of segments that were received out of order at the
    MPTCP (meta) level."
 ::= { mptcp 19 }
```

mptcpSentSegments OBJECT-TYPE

```
SYNTAX      Counter
UNITS       "segments"
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The number of segments that were emitted at the
    MPTCP (meta) level."
 ::= { mptcp 20 }
```

END

3. IANA considerations

The MIB module in this document uses the following IANA-assigned OBJECT IDENTIFIER values recorded in the SMI Numbers registry:

Descriptor	OBJECT IDENTIFIER value
----- mptcp	----- { mib-2 TBD}

RFC Editor: The IANA is requested to assign a value for "TBD" under the 'mib-2' subtree and to record the assignment in the SMI Numbers registry. When the assignment has been made, the RFC Editor is asked to replace "TBD" (here and in the MIB module) with the assigned value and to remove this note.

4. Conclusion

This document has proposed a simple Management Information Base (MIB) to manage Multipath TCP.

5. Acknowledgements

This work was partially supported by the FP7 TrilogY-2 project funded by the EC.

6. Normative References

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

Authors' Addresses

Fabien Duchene
UCLouvain

Email: fabien.duchene@uclouvain.be

Christoph Paasch
Apple

Email: cpaasch@apple.com

Olivier Bonaventure
UCLouvain

Email: Olivier.Bonaventure@uclouvain.be