                Considerations for MPTCP operation in 5G
            draft-purkayastha-mptcp-considerations-for-nextgen-00

Abstract

   MPTCP is deployed in devices which have multiple interfaces to
   different networks. It takes advantage of these multiple interfaces
   to improve user experience by bandwidth aggregation, redundancy, and
   increased reliability by having a fallback option. This document
   describes scenarios where next generation (5G) mobility management
   frameworks may impact MPTCP operations thus potentially jeopardizing
   achievement of some of these gains.

Table of Contents

1. Introduction

Utilization of MPTCP in devices with multiple interfaces is very
attractive. MPTCP is being deployed and widely adopted in today's
smart devices, which typically have multiple network interfaces such
as Cellular and Wifi. It provides reliability, bandwidth aggregation
capability, and handover efficiency. MPTCP assumes that an interface
is associated to only one IP address. This document describes next
generation mobility management frameworks that may assign more than
one IP address (to a given user device) over a single interface. In
this scenario, MPTCP may not produce desired results.

This document describes the anticipated difficulties for MPTCP when
working on devices which are being managed by modern mobility
management protocols. The goal of this document is to make the WG
aware of the upcoming situation and gauge interest in working to
preserve MPTCP efficiency.

2. MPTCP background

Multipath TCP is an extension to TCP allowing hosts to use multiple
paths to send and receive the data belonging to one connection. The
simultaneous use of these multiple paths (sub-flows) for a TCP/IP
session improves resource usage within the network and, thus,
improves user experience through higher throughput and improved
resilience e.g. to network failure.

A Multipath TCP connection is composed of several TCP connections
that are called sub-flows. MPTCP manages the creation, removal, and
utilization of these sub-flows to send data.

Once an MPTCP connection has begun, further sub-flows can be added to
the connection. Hosts have knowledge of their own addresses, and can
become aware of the other host's addresses through signaling
exchanges. If extra paths are available, additional TCP sessions
("sub-flows") are created on these paths, and are combined with the
existing MPTCP session, which continues to appear as a single
connection to the applications at both ends. Standard TCP versus
MPTCP protocol stacks are illustrated in Figure 1, from [RFC6824].

```
                            +---------------------------------+
                            |           Application           |
        +--------------+    +---------------------------------+
        | Application  |    |            MPTCP                 |
        +--------------+    +---------------------------------+
        |    TCP       |    | Subflow(TCP)  |  Subflow(TCP)  |
        +--------------+    +---------------------------------+
        |    IP        |    |    IP         |     IP          |
        +--------------+    +---------------------------------+
```
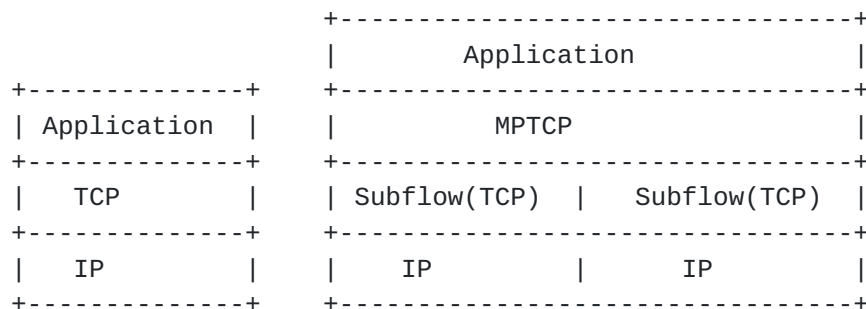
Figure 1: Standard TCP vs MPTCP Protocol Stacks

3. Mobility management protocols

Mobility management is needed because the IP address of a mobile node
may change as the node moves. A fair number of network-layer (IP)
mobility protocols have been standardized. They all employ a
mobility anchor to allow a mobile node to remain reachable after it
has moved to a different network. 3GPP 5G has adopted distributed
approach and pushed the mobility anchor towards the edge of the
network.

3.1. Centralized mobility management (e.g. DSMIP, PMIP)

Existing network-layer mobility management protocols have primarily
employed a mobility anchor to ensure connectivity of a mobile node by
forwarding packets destined to, or sent from, the mobile node after
the node has moved to a different network. The mobility anchor has
been centrally deployed in the sense that the traffic of millions of
mobile nodes in an operator network is typically managed by the same
anchor. Redirecting packets this way can result in long routes.

3.2. Mobility Handling in 5G

   3GPP standards organization is currently defining 5G network
   architecture and systems. This is referred to as multi-homed PDU
   Session.

   In 5G networks, mobility framework under development suggest that IP
   mobility, network access solutions, and forwarding solutions should
   enable traffic to avoid traversing a single mobility anchor far from
   the optimal route. Mobility anchors are distributed in such a way
   that the mobility anchors are positioned closer to the user; ideally,
   mobility agents could be collocated with the first-hop router.

   The current definition specifies that a PDU (packet data unit)
   Session may be associated with multiple IPv6 prefixes. The PDU
   Session provides access to the Data Network via more than one PDU
   (IPv6) anchor. The different user plane paths leading to the IP
   anchors branch out at a "common" UPF (user plane function) referred
   to as a UPF supporting "Branching Point" functionality. The
   Branching Point (BP) provides forwarding of uplink traffic towards
   the different IP anchors and merge of downlink traffic to the mobile
   device i.e. merging the traffic from the different IPv6 anchors on
   the link towards the device. The BP and UPFs (IP anchors) are
   illustrated on Figure 2, as described in [3GPP.23.501].

```
                          +--------+    +-------+
              --------------| AMF    |----| SMF   |------
             |              +--------+    +-------+     |
             |          |                     |        +-----------+  +----+
             |          |                     |      /-|  UPF       | |    |
    +-------+    +--------+    +-----------+  /  |PDU Session|--|    |
    |       |   |        |    |           | / |Anchor 1   |  |    |
    | UE    |---|  AN    |----| Branching |/    +-----------+  |    |
    |       |   |        |    | Point     |\                    | DN |
    +-------+    +--------+    +-----------+ \   +-----------+  |    |
                                              \--|  UPF       | |    |
                                                 |PDU Session|--|    |
                                                 |Anchor 2   |  |    |
                                                 +-----------+  +----+
```
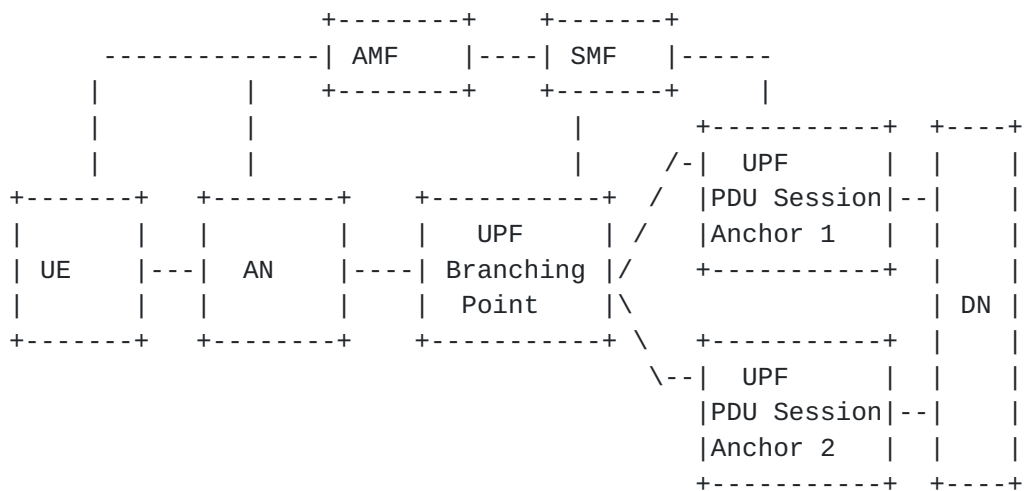
             Figure 2: Multi-Homed PDU Session: Service Continuity Case

4. Considerations for MPTCP

   Using MPTCP (as currently defined) in 5G network, which may assign
   multiple IP addresses (to a given device) on the same interface, may
   need special consideration.

   As we know, MPTCP currently interprets multiple IP addresses as
   indicating separate distinct network interfaces. This incorrect
   interpretation may lead to not realizing fully the benefit of MPTCP
   applications such as bandwidth aggregation, redundancy, and
   reliability in 5G systems.

5. MPTCP operation use case

5.1. BW management increase BW as needed and load balancing features

   One of the main goal of using MPTCP is to increase the BW associated
   to a TCP session when needed. The additional BW is obtained by
   creating sub-flows over interfaces other than the one already in use
   (e.g. using different IP addresses) and associating these sub-flows
   to the existing TCP session, which altogether form the MPTCP session.
   The total BW available for an MPTCP session is the accumulation of
   all the sub-flows associated to the same MPTCP session.

   In situations, when MPTCP sub-flows are created using IP addresses
   associated to the same interface, MPTCP operation is impacted.
   Adding up sub-flow using the same interface doesn't increase the
   available BW (considering that the bottleneck may be between the
   device and the PoA). In this case, MPTCP may continue to add more
   sub-flows but there may not be proportional gain in bandwidth. In
   this context, MPTCP functionality is inefficient; the expected MPTCP
   behavior not being met.

   Similar situation arises for load balancing. MPTCP may try to add
   sub flows associated to the same interface and attempt to distribute
   the load. From operational perspective, we may not see proportional
   gain in load balancing.

5.2. Backup path management - created in advance to be available when
      needed

   MPTCP may associate sub-flows to an MPTCP session with one or many of
   these sub-flows being identified as "backup" sub-flows. Backup sub-
   flows are used only when no regular sub-flows are available (e.g. if
   a link failure occurs).

   As described in the previous section, the backup mechanism is
   ineffective if the backup path is using the same interface as the

regular path. If the interface goes down, then the backup path
becomes unavailable, at the same time as the regular path. In this
case, no backup path is ready to handle data traffic. New sub-flows
will need to be created to backup the regular sub-flow, which is
unavailable. Packets loss and retransmissions will become inevitable
resulting in bad user experience.

6. IANA Considerations

This document requests no IANA actions.

7. Security Considerations

TBD

8. Informative References

[_3GPP.23.501]
           3GPP, "System Architecture for the 5G System", 3GPP
           TS 23.501 1.4.0, 9 2017,
           <http://www.3gpp.org/ftp/Specs/html-info/23501.htm>.

[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>.

Authors' Addresses

   Debashish Purkayastha
   InterDigital Communications, LLC
   Conshohocken
   USA

   Email: Debashish.Purkayastha@InterDigital.com

   Michelle Perras
   InterDigital Communications, LLC
   Montreal
   Canada

   Email: Michelle.Perras@InterDigital.com

   Akbar Rahman
   InterDigital Communications, LLC
   Montreal
   Canada

   Email: Akbar.Rahman@InterDigital.com