MPTCP Working Group   L. Deng
INTERNET-DRAFT   D. Liu
Intended Status: Informational   T. Sun
Expires: May 25, 2014   China Mobile
                                                         M. Boucadair
                                                       France Telecom
                                                           G. Cauchie
                                                     Bouygues Telecom
                                                         Oct 24, 2014

Use-cases and Requirements for MPTCP Proxy in ISP Networks
draft-deng-mptcp-proxy-01

Abstract

This document presents the use-cases and identifies requirements for ISP deployed MPTCP proxies for both Fixed and Mobile networks.

Status of this Memo

Copyright and License Notice

Table of Contents

1 Introduction

Internet Service Providers (ISPs) are challenged by the data explosion occurring in their Fixed and/or Mobile networks. This data explosion is triggered by new usages with the advent for resource-demanding services. The emergence of these services is facilitated by the emergence of new access technologies (FTTx in the Fixed or LTE in the Mobile networks). Typical resource-demanding services are (HD) video streaming or catchup IP-TV which are boosting more and more every day the customers appetite for more IP bandwidth.

This pressure on continuous increase of network capacity poses a challenge on the ISPs to appropriately plan, dimension and dynamically provision their networks to satisfy customers expectations. This problem is encountered by both Fixed and Mobile Providers that have to cope with the scarcity of the radio frequency resources, the limits of the already widely deployed DSL infrastructure, or any in-place access technology.

The traditional trend that consist in upgrading the access technology to a new generation technology may show some limits because of the required upgrade cycles. Alternate deployment options to satisfy customers' expectations and react rapidly to competition are of interest of ISPs. A promising track, discussed in this document, is to aggregate several connectivity links while eliminating risks to experience application failures.

Indeed, a direction to answer to this problem is to make use of the multiple interfaces that one terminal host maintains. For smartphones or mobile dongles, these interfaces are typically a 3G/4G radio access and a WLAN access, while for a residential CPE these interfaces are typically a DSL line and a 3G/4G radio access.

3GPP initiated an effort to aggregate several radio resources for the sake of increasing bit-rates (denoted as Carrier Aggregation (CA)). Aggregation is achieved at the radio level by combining the set of allocated contiguous or non-contiguous component carriers. This extension requires modification at the radio interface level of the UE (User Equipment), as well as some tuning on the network side. Carrier Aggregation is specific to radio-based environments and, as such, it is not convenient for other deployment cases, such as wired networks.

Both 3GPP and IETF standard bodies have investigated different solutions to make the best interface used for one application, with potential use of multiple interfaces at the same time by multitasking-enable terminals. As of today, none of them really did meet a wide deployment and successful adoption.

The IETF released MPTCP specification [RFC6824], an experimental set of TCP options allowing the use of parallel TCP connections, each with different set of IP addresses and/or port numbers, to serve at least one application. MPTCP is already available on some widely adopted mobile handsets and on some Linux implementations. However, MPTCP connections are pretty rare since the servers hosting the applications are too few to offer MPTCP capability.

Because resources are scarce at the access segment, a solution to enhance the quality of experience is to enable MPTCP at the access segment without requiring MPTCP support at the server side (i.e., the end-to-end MPTCP support is not required). Concretely, this can be implemented owing to the deployment a dedicated function called MPTCP Proxy at the ISP network side and/or in the CPE device. Note that enabling an MPTCP Proxy in the CPE has the advantage to not require MPTCP stack at the terminal side.

By this mean, an ISP can offer a higher bandwidth to its customers when possible while not waiting for massive MPTCP adoption by the Internet ecosystem. Furthermore, this technique would allow the ISP and the customer to control the parts of the networks where potential QoE degradation may be experienced and where the traffic can be handled appropriately by means of traffic engineering tweaking for instance. Since MPTCP requires a load-sharing algorithm to schedule the TCP subflow to which the traffic is forwarded, the selection of the proper algorithm could help for instance to offload the IP traffic towards the legacy Fixed networks while taking advantage of the complementary bandwidth only when needed/selected by the customers, application, or the ISP.

Note that the use of MPTCP at customer side can be of different natures as defined in MTPCP base specification:

o Native MPTCP: Two MPTCP endpoints establish and make use of all subflows that correspond to the available addresses/port numbers. This mode is enabled to optimize data throughput.

o Active/Backup MPTCP: Two MPTCP endpoints enable multiple subflows, but only a subset of these subflows is actually in use for data transfer. MPTCP endpoints can use the MP_PRIO signal to change the priority for each subflow.

o Single subflow MPTCP: Two MPTCP endpoints use one single subflow and when a failure is observed, an additional subflow is enabled so that traffic is forwarded along the newly established subflow.

Based on the basic capabilities of an MPTCP-enabled host, various deployment use cases can be considered by an ISP. Examples of such

use cases include: traffic handover among multiple WLAN hotspots, traffic offload from a mobile network to WLAN or vice visa, and access link aggregation.

In general, two flavors of MPTCP Proxy can be envisaged , see Figure 1:

o A MPTCP Proxy that maps a UE originated TCP connection into an MPTCP connection. An example would be an MPTCP-enabled CPE, which makes use of its own multiple local interfaces to maximize the experience of bandwidth consuming applications for Non-MPTCP UEs.

o A MPTCP Proxy that maps an UE originated MPTCP connection into a TCP connection. This is the case for most mobile terminals with multiple interfaces and built-in MPTCP capability.

```
+----+        +------------+         +------------+       +------+
|Host| <=TCP=>| MPTCP Proxy|<=MPTCP=>| MPTCP Proxy|<=TCP=>|Server|
+----+        +------------+         +------------+       +------+

+----+         +------------+       +------+
|Host|<=MPTCP=>| MPTCP Proxy|<=TCP=>|Server|
+----+         +------------+       +------+
```

Figure 1: MPTCP Behaviors.

This documents details these use-cases and derived requirements for MPTCP proxy deployment in ISP networks.

The use cases discussed in this document can also be valid for Enterprise networks.

2 Terminology

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].

This document makes use of the following terms: (For definitions of other terms such as "(MPTCP) Connection", "Host" and "Subflow", refer to [RFC6824].)

M-session: a MPTCP connection between a M-UE and a M-Server.

M-Server: a server with MPTCP capability enabled for the current TCP

session, or simply a serving MPTCP host. M-Server may be connected to the network using one or multiple network interfaces.

M-UE: a user equipment with MPTCP capability enabled for the current TCP session, or simply a requesting MPTCP host. Unless it is explicated, an M-UE can be a host or a CPE.

N-Server: a server without MPTCP capability enabled for the current TCP session.

N-UE: a user equipment without MP-TCP capability enabled for the current TCP session.

MPTCP Proxy: proxy located between a M-UE and a N-Server, which enables a M-session with the M-UE and maps it into a legacy TCP connection with the N-Server.

Natural Path: a "natural path" of a subflow is the original path it would traverse end-to-end if there is no explicit traffic steering function is enabled for MPTCP Proxy.

Small Cell: generally refers to the cells with less than 1 watt output power, e.g. picocell, femtocell, etc.

Nanocell: a type of base station integrated with both a small cell and a WLAN access point.

3 Use-cases for ISP MPTCP Proxy

3.1 Boosting MPTCP Utilization for M-UEs and Multi-Interface CPEs

On the one hand, user equipment (esp. mobile terminals) nowadays have multiple interfaces and could benefit from these interfaces if they were using MPTCP.

On the other hand, servers often only support regular TCP and upgrading them to support MPTCP will take more time than on the clients.

Therefore, it is expected that an ISP deployed MPTCP Proxy would benefit the M-UEs with better user-experience for almost all the legacy applications by exploiting MPTCP capability at least for the most resource-limited channel on the air, without relying on pervasive MPTCP deployment at the server side.

The same approach can be followed for CPE-based deployment models. These models can be refreshed to integrate CPEs that are able to

mount multiple interfaces.

3.2 Resource Pooling from Multiple Networks

For an ISP providing multiple access networks to its subscribers, a locally deployed MPTCP Proxy would enable use-cases where the M-UE/ISP is trying to pool access network resources from multiple networks concurrently, e.g., among multiple WLAN hotspots, a fixed and a WLAN connection, a cellular and a WLAN, a cellular and a Fixed access, etc.

However, the specific pooling strategy can differ for various scenarios, depending on the type of subscriber and the current status of different networks, etc.

For example, from the user's perspective, business customers would have perfect pooling while other users would get the cheaper network whenever possible.

On the other hand, from the ISP's perspective, it would encourage more efficient usage of network resource by dynamic charging policy for rush hours, has a static preference for a specific network over another one, or even desires traffic offloading to proactively migrate less sensitive or more demanding traffic between multiple networks it controls.

Criteria to invoke an MPTCP Proxy in a communication paths will be the results of these policies (i.e., subscribers, applications, and ISPs).

3.3 Multiple Connections and Seamless Handover between Multiple Networks

For cellular ISPs, the radio access networks are the dominating part of their network construction investment. With the rapid development of cellular technology and the imbalance of subscriber/traffic distribution geographically, it is common practice to deploy leading-edge technology in traffic boosting hot spots (e.g., downtown areas in big cities) first while enabling seamless handover for legacy cellular/wireless networks to guarantee service continuity for a moving terminal.

A device can discover multiple networks, including multiple attachment points to the same ISPs. This is the typical example of a device that can be serviced with multiple WLAN hotspots. Instead of selecting only one attachment point, the device can select multiple ones when more resources are required. Maintaining simultaneous attachment to these attachment points will also allow for seamless handover and, therefore, allow for session continuity.

Implementing this feature may require an explicit signal to the connecting device to drive its network attachment procedure. It is out of scope of this document to discuss such signals nor elaborate on potential impact on battery consumption on mobile devices mounting multiple interfaces in parallel.

3.4 Assist MTPCP Connection Establishment

The MPTCP Proxy does not know in advance whether a remote server is MPTCP-enabled. As such, the MPTCP Proxy can be provided with various policies such as (but not limited to):

o TAKE-OVER Mode: Strip any MPTCP signal systematically in all TCP messages to a remote server: This mode might be useful for two scenarios: For the scenario where servers are not widely MPTCP-enabled, it has the advantage to not suffer from MPTCP fallback delay when the remote server is not MPTCP-enabled. For the scenario where the ISP would like to limit the MPTCP traffic to its local network to avoid unexpected blocking by third-party middle-boxes.

o TRANSPARENT Mode: Allow MPTCP signals to be passed to systematically to remote servers: This mode has the advantage to optimize MPTCP Proxy resources and favor direct communications between an MPTCP-enabled UE and an MTPCP-enabled server without invoking the MPTCP Proxy when both the server and client are MPTCP-enabled.

o HYBRID Mode: The combination of the above two modes, according to the proxy's local policy.

4 Deployment Considerations

For an MPTCP Proxy to correctly manage an M-session with the M-UE, it is necessary for it to receive each subflow coming from the M-UE and heading for the N-Server it is acting on behalf of. It is hence straightforward to consider an on-path deployment strategy, where the MPTCP Proxy is directly located on the common link of every subflow from the M-UE.

However, for other cases where a common link is absent within the local ISP's domain or resource pooling is desired from networks of different ISP domains, the MPTCP Proxy has to steer off-path subflow to traverse the selected MPTCP Proxy explicitly. Note that steering functions should work for both directions over all the subflows, i.e. including both uplink and downlink traffic from/to the M-UE.

Therefore, two types of MPTCP Proxy are considered in this document: on-path MPTCP Proxy and off-path MPTCP Proxy.

4.1 On-path MPTCP Proxy

For different access networks provided by a single ISP, it is fairly easy to locate a common link and deploy an on-path MPTCP Proxy accordingly.

As depicted in Figure 2, the on-path MPTCP Proxy is located on each potential path from a M-UE for both MPTCP traffic and legacy TCP traffic.

```
                    (           )
                 ( Access Net #1  )
         +--->(e.g. Cellular Network)<--+
+----+   | ....>(                 )<.   |  +-------+    +--------+
|    |<--+ .        (           )   .   +->|       |<==>|        |
|    |<.....                        ......>|On-Path|<..>|        |
|M-UE|                              +----->| MPTCP |    |N-Server|
|    |<----+                        | ....>| Proxy |<..>|        |
|    |<... |       (           )    | .    +-------+    +--------+
+----+   . +--->(  Access Net #2  )<+ .
         ....>( e.g. WLAN Network   )<.
                (                 )
                   (           )


     -----subflow of an M-session
     .....legacy TCP flow
     =====merged TCP flow for an M-session
```

Figure 2. The On-Path MPTCP Proxy

4.2 Off-Path Proxy

On the one hand, it is not viable to assume for instance a common link from a cellular ISP to deploy an on-path proxy on each potential MPTCP subflow originating from its M-UE via third-party WLAN networks for instance.

On the other hand, even for resource pooling from a single ISP's own WLAN networks, it is not always feasible to find such a common link for on-path MPTCP Proxy before corresponding traffic hitting the public Internet.

```
                    (           )
```

```
                    ( Access Net #1  )
            +--->(e.g. Cellular Network)<--+
   +----+   | ....>(                 )<.   |  +--------+    +--------+
   |    |<--+ .       (           )    .   +->|off-Path|<==>|        |
   |    |<.....                         ......>|  MPTCP |<..>|        |
   |M-UE|                               ******>|  Proxy |    |N-Server|
   |    |<*****                         *      +--------+    |        |
   |    |<... *       (           )    * ..................>|        |
   +----+   . *****>(  Access Net #2 )<* .                  +--------+
            ....>(  e.g WLAN Network   )<.
                   (                 )
                      (           )

     -----subflow following its natural path
     *****explicitly redirected subflow
     .....legacy TCP flow
     =====merged TCP flow for a M-session
```

Figure 3. The Off-Path MPTCP Proxy

As depicted in Figure 3, the off-path MPTCP Proxy is located on the "natural path" of only a partial subset of potential subflow of a given M-session, and relies on extra mechanism to explicitly steering other subflows to deviate from their "natural path" and go through it.

5 Deployment Scenarios

5.1 MPTCP Proxy at the ISP IP Gateway

For Fixed networks, MPTCP Proxy can be located in various segments with a network (e.g., co-located with a CGN [RFC6888], a DS-Lite AFTR [RFC6333], or a NAT64 device [RFC6146] if already present in the network, co-located with a TCP acceleration and optimizer if any, etc.).

As the anchoring point for 3GPP mobile UE originated IP traffic within the cellular network, the IP gateway of 3GPP network (e.g., GGSN (Gateway GPRS Support Node) or PDN Gateway (PGW) [RFC7066]) would be a natural spot for MPTCP Proxy deployment. It is also possible for the gateway to reuse existing interfaces with other 3GPP network elements for information/policy acquisition. Alternatively, another location for MPTCP Proxy deployment would be the (s)Gi Interface(i.e., between the GGSN/PGW and an external PDN (Packet Domain Network)).

5.2 On-Path MPTCP Proxy at the ISP Data Center

It is common practice for local ISPs to build up local data center facilities within its domain for large Internet content providers in order to feed user's requests locally, resulting in both enhanced user experience for cut-short end-to-end delay and reduced expenses for unnecessary cross-ISP peering traffic.

It is believed that by deploying an on-path MPTCP Proxy at the entrance of the ISP's local data center, it would help various Internet Content Provider residing within the data center to gain enhanced user-experience for local subscribers with M-UEs without the need to upgrade their servers manually.

5.3 On-Path MPTCP Proxy at SmallCell Gateway

Some ISPs are deploying small cells (low-powered radio access nodes) with both cellular and carrier WLAN access. Small cells is expected to be a cost-effective and green way for cellular network deployment for both home and enterprise subscribers. For instance, it can be integrated with the home gateway for a broadband subscriber.

Figure 4 outlines the Nanocell system architecture.

```
          (    )
+--+  +-(Cellular)-+ +--------+                 Cellular
|  |-+    (    )   +-|Nanocell|              +-(  Core   )---(    )
|UE|                 |(local  |-(IP Backhaul)+             (Internet)
|  |-+    (    )   +-|Gateway)|              +-( WLAN AC )---(    )
+--+  +-(  WLAN  )-+ +--------+
          (    )
```

Figure 4. Nanocell System Architecture

As a combined access node for both cellular and WLAN traffic, small cell Gateway represents a spot for an on-path MPTCP Proxy at the network edge.

Note that by applying local breakout scheme, where the small cell doing IP-layer forwarding itself rather than relying on other 3GPP routing devices, it is expected that no extensions to existing 3GPP specs are needed for its integration with an MPTCP Proxy.

5.4 MPTCP Proxy at CPE/CGN for Fixed Access Networks

A first deployment model assumes MPTCP Proxy is enabled in a CPE. This proxy aims to offer MPTCP service to non-MPTCP hosts located

behind the CPE. This model can be deployed with or without an MPTCP Proxy at the ISP's network.

In early deployment stages, this model is likely to require an MPTCP Proxy be also deployed at the ISP's network side. This is mainly to shorten delays induced by TCP fall back when the remote server is not MPTCP compliant.

Hosts located behind the CPE are not aware of the activation of the MPTCP at the CPE side. MPTCP can be used for both bandwidth aggregation and also ensure session continuity during failure events.

6 Requirements for MPTCP Proxy

This section identifies requirements derived from the above use-cases for an ISP MPTCP Proxy.

6.1 Protocol Proxying between MPTCP and TCP

In the on-path MPTCP Proxy use-case, the proxy is assured to be deployed on the path of each potential subflow of a given MPTCP session from the M-UE to the N-Server.

To allow a MPCTP-enabled UE to make full use of the multiple interfaces even if it is interacting with an N-server, the on-path MPTCP Proxy MUST satisfy the following requirements.

(a) Compatibility: An on-path MPTCP Proxy supports detection of M-UE/N-server combinations for further proxying while leaving M-UE/M-server and N-UE/N-server sessions intact.

(b) Transparency: An on-path MPTCP Proxy supports negotiation with and acting towards the M-UE like a M-server on behalf of N-Server, while acting towards the N-Server like a N-UE on behalf of the M-UE.

6.2 Explicit Traffic Steering for Off-Path Proxying

As we discussed earlier in the off-path MPTCP Proxy use-case, it is required that subflows whose "natural path" does not include the MPTCP Proxy be redirected explicitly to go through the proxy.

(c) Explicit Traffic Steering: an off-path MPTCP Proxy that enables resource pooling from third-party WLAN networks MUST support explicit traffic steering, to allow all the subsequent subflow traffic go through the exactly the same MPTCP Proxy used in the correspondent M-session establishment for both directions (including uplink and

downlink traffic from/to the M-UE).

(d) Globally Routable Address: an off-path MPTCP Proxy that enables resource pooling from the same ISP's networks SHOULD expose a globally routable address to allow explicit steering of subsequent subflow traffic after they hit the public Internet.

6.3 Flexible Resource Policy within a Single ISP

Different from the end-to-end MPTCP solution, ISP deployed MPTCP Proxy is a potential point for centralized cross-network flow control over service/RAT preference for a given subscriber's M-UEs, which is considered to be essential for better operation and service provision in 3GPP networks.

However, to enable such fine-grained resource pooling policy from the network side, it is essential that the MPTCP proxy knows for each subflow its specific Network Access Type information.

(e) Network Access Type Information: an MPTCP proxy SHOULD be able to make use of a reliable information sharing/reporting mechanism to acquire a subflow's Network Access Type information/update on a real-time basis.

(f) Resource Policy: an MPTCP Proxy MUST support flexible control to set limits to both the number of concurrent subflows running in a M-session and the number of concurrent M-sessions from an M-UE/to an N-Server, according to the type and/or identity of relevant subscriber and/or application.

6.4 Protection against third-party traffic

Apart from the traditional private network deployment practice, an off-path MPTCP proxy exposes itself as publicly accessible from any third-party traffic, where traditional access authorization or admission control mechanisms from 3GPP network would not work. It is therefore envisioned that an off-path MPTCP Proxy open for third-party WiFi resource pooling MUST support minimal protection/policy against potentially malicious traffic from third-party network.

(g) Provision Negotiation: an MPTCP Proxy SHOULD support both subscriber/M-session/subflow level resource reservation negotiation with a M-UE.

(h) Origin Authentication: an off-path MPTCP Proxy MUST support subflow authentication for traffic from an unauthorized third-party WiFi, in order to serve subflows belong to its intended M-sessions coming from authorized subscribers or NAT, while turning down others

with least overhead.

(i) Admission Control: an off-path MPTCP Proxy MUST support admission control to set limits to both the number of concurrent subflows on a given M-session and the number of concurrent M-sessions for a given subscriber/application.

6.5 MPTCP Proxy Selection from Multiple Candidates

When multiple MPTCP Proxies exist on an end-to-end path from the M-UE to the N-Server, a natural question arises that "which MPTCP Proxy should be chosen and how". There may be different considerations depending on the location and types of MPTCP Proxy involved in addition to the expectation of the application.

For example, assume an ISP deploys on-path MPTCP Proxy at smallcell Gateway as well as an off-path MPTCP Proxy at the IP gateway of its cellular network. The following considerations may apply.

On one hand, a straightforward policy would be to choose the nearest MPTCP Proxy to the M-UE, i.e. the on-path MPTCP Proxy at smallcell gateway which would lead to more efficiency from less traffic convergence pressure. Another policy would be to prefer the on-path MPTCP Proxy to the off-path one, which would cause unnecessary traffic roundabout.

However, on the other hand, if a M-session is established with the on-path MPTCP Proxy at smallcell gateway and when the M-UE moves out of the coverage of the cell, the connection will be broken. Therefore, it is reasonable to choose off-path MPTCP Proxy for applications with strict service continuity expectations, while favor on-path MPTCP Proxy for bulk data transfer with application-level re-connection mechanisms.

In summary, regarding the MPTCP Proxy selection from multiple candidates, the following requirement apply.

(j) Flexible Selection: when multiple M-Proxies are deployed on the end-to-end path for a M-session establishment within the domain a single ISP, it SHOULD be possible for the ISP to enforce flexible selection policy regarding which MPTCP Proxy to serve which M-session, based on the MPTCP Proxy's location, the MPTCP Proxy's type (on-path/off-path) in addition to the application's expectation.

6.6 Load Balancing Algorithm for Multiple Networks

When multiple networks are available to service a subscriber, the traffic may be balanced among those available paths for the sake of

QoE. The logic for balancing the traffic among those multiple paths can be driven by application needs, customer's preferences, and/or ISP traffic engineering guidelines.

From a traffic engineering perspective, the ISP may enforce policies that would optimize various parameters such as:

o Network resources usage as a whole.

o Optimized invocation of available MPTCP Proxies (i.e., MPTCP Proxy selection).

o Optimized MPTCP Proxy local performances (e.g., avoid overload phenomena).

o Enhanced QoE (including increase both upstream and downstream throughputs)

In summary, regarding the load balancing algorithm for multiple networks, the following requirement apply.

(k) The MPTCP Proxy SHOULD be configurable with the load balancing ratio per each available path.

6.7 Misc

Below are listed some additional requirements:

o Including an MPTCP Proxy in the communication may be seen as a single point of failure. Means to protect against such failures should be enabled.

o If TCP flows handled by the MPTCP Proxy are sourced with the external IP address(es) that belong to the MPTCP Proxy, all hosts serviced by the same MPTCP Proxy will be seen with the same IP address(es). Means to mitigate the side effects documented in [RFC6269] SHOULD be enabled in such deployment case.

o The MPTCP Proxy SHOULD NOT alter non-MPTCP signals included in a TCP segment.

o The MPTCP Proxy MUST NOT inject MPTCP signals if the TCP option size is consumed.

o The MPTCP Proxy SHOULD NOT inject MPTCP signals if this leads to local fragmentation. MTU tuning may be required to avoid fragmentation. MPTCP Proxy SHOULD be configurable to enable/disable this feature.

- o The MPTCP Proxy MUST take proper measure to avoid errors caused by careless MPTCP signal modification to segments with other TCP options. For instance, TCP-AO (TCP Authentication Option) [RFC5925], which includes TCP options in the MAC computation, when present, MUST be the first processed by the MPTCP Proxy.

- o The MPTCP Proxy SHOULD be easy to scale to cope with growing demand.

7 Security Considerations

As an MPTCP Proxy is playing N-UE and M-Server at the same time, the security considerations that concerns a TCP client and a MPTCP-enabled server are applicable to a MPTCP Proxy in general [RFC6181].

Forcing the traffic to cross an MPTCP Proxy that would not be involved if legacy routing and forwarding actions are enforced raises some security concerns. In particular, inserting an illegitimate MPTCP Proxy can be used to hijack connections. Traffic inspected by third party MPTCP Proxy may be used as a vector to reveal privacy-related information.

These security concerns can be mitigated if the MPTCP Proxy is managed by the same ISP offering connectivity to a customer. Otherwise, specific mechanisms to be used are expected to be an integral part of an MPTCP Proxy design and out of scope of this document.

8 IANA Considerations

There is no IANA action in this document.

9 Acknowledgement

The authors would like to thank Olivier Bonaventure, Jordan Melzer, Preethi Natarajan, Marc Portoles, Chunshan Xiong, Alper Yegin, Zhen Cao and Hui Deng for their valuable comments and input to this document.

10 References

10.1 Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

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

10.2 Informative References

[RFC6181] M. Bagnulo, "Threat Analysis for TCP Extensions for Multipath Operation", RFC 6181, March 2011.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC6333, August 2011.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

[RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, June 2011.

[RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, June 2011.

[RFC7066] Korhonen, J., Arkko, J., Savolainen, T., and S. Krishnan, "IPv6 for Third Generation Partnership Project (3GPP) Cellular Hosts", RFC 7066, November 2013.

[RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, June 2010.

Authors' Addresses

Lingli Deng
China Mobile

Email: denglingli@chinamobile.com

Dapeng Liu
China Mobile

Email: liudapeng@chinamobile.com

Tao Sun
China Mobile

Email: suntao@chinamobile.com

Mohamed Boucadair
France Telecom

Email: mohamed.boucadair@orange.com

Gregory Cauchie
Bouygues Telecom

Email: GCAUCHIE@bouyguestelecom.fr