

QUIC Working Group
Internet-Draft
Intended status: Informational
Expires: December 1, 2020

M. Boucadair, Ed.
Orange
O. Bonaventure, Ed.
M. Piraux, Ed.
Q. De Coninck
UCLouvain
S. Dawkins, Ed.
Tencent America
M. Kuehlewind, Ed.
Ericsson
M. Amend
Deutsche Telekom
A. Kassler
Karlstad University
Q. An
Alibaba Group
N. Keukeleire
Tessares
S. Seo
Korea Telecom
May 30, 2020

3GPP Access Traffic Steering Switching and Splitting (ATSSS) - Overview
for IETF Participants
[draft-bonaventure-quic-atsss-overview-00](#)

Abstract

This document briefly presents the Access Traffic Steering, Switching, and Splitting (ATSSS) service being specified within the 3rd Generation Partnership Project (3GPP). The ATSSS service provides network support for multihomed devices to select a path for transmission (steer), move traffic from one path to another (switch), or use multiple paths simultaneously (split). TS 23.501 specifies an ATSSS architecture for TCP traffic.

This document presents a snap-shot of the ongoing discussion in the 3GPP to enable ATSSS for non-TCP traffic, based on the use of QUIC, and assesses to what extent IETF specifications can be used to meet the ATSSS design goals. Apparent gaps are also documented.

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 December 1, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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	3
1.1.	Notes for Readers	4
2.	Introduction to Access Traffic Steering, Switching, and Splitting (ATSSS)	4
3.	Contribution and Discussion Venues for this draft.	5
4.	Conventions, Terminology, and Definitions	6
5.	High Level ATSSS Overview	7
5.1.	Reference Architecture	7
5.2.	External IP Addresses Used by the ATSSS UPF	8
5.3.	ATSSS Modes	9
5.4.	ATSSS Rules	9
6.	ATSSS Phases	11
7.	ATSSS Phase 1: Support for TCP	11
7.1.	ATSSS Phase 2: Adding Non-TCP Support	13
7.1.1.	QUIC and Multihoming	14
7.1.2.	QUIC as an ATSSS Data Plane Protocol	15
7.1.3.	Single QUICv1 Tunnel with Unreliable Datagram Extension and Connection Migration	15
7.1.4.	Multiple QUICv1 Tunnels with Unreliable Datagram	

Extension and Connection Migration	17
7.1.5. MP-QUIC Tunneling	18
7.2. Mapping of Both TCP and Non-TCP to QUIC Streams and Datagrams	19
7.3. Encapsulation Overhead	20
7.4. Multiple Encryptions	21
7.5. Congestion Control in Congestion Control and Coexistence	21
7.6. Packet Order Reconstruction for (MP-)QUIC Splitting Mode	22
8. QUICv1 Gap Analysis for ATSSS Phase 2	22
9. IANA Considerations	23
10. Security Considerations	23
11. Acknowledgements	24
12. Document History	24
13. References	24
13.1. Normative References	24
13.2. Informative References	24
Authors' Addresses	27

[1.](#) Introduction

The 3rd Generation Partnership Project (3GPP) has described the Access Traffic Steering, Switching, and Splitting (ATSSS) service, used to carry traffic over multiple available paths. In Release 16 [[TS23501](#)], ATSSS supports TCP traffic relying upon two IETF protocols: MPTCP [[RFC6824](#)] and Convert Protocol [[I-D.ietf-tcpm-converters](#)].

As part of preparation for Release 17 studies, 3GPP has expressed an interest in other IETF protocols and protocol extensions that would enable ATSSS service of traffic not supported by the Convert Protocol nor based on the use of MPTCP. To that aim, 3GPP has contacted the IETF through a formal liaison [[atsssliaison](#)], letting the IETF know about this interest. An excerpt of the liaison document is provided below:

"The work on the study has not yet started in 3GPP, and there are thus no agreed conclusions. The goal is to enable steering, switching and splitting of traffic (primarily UDP) across multiple accesses, including latency sensitive and real time traffic. Therefore 3GPP is interested to receive regular feedback on progress and prioritization on the multipath extensions to QUIC."

Because "3GPP SA kindly requests IETF to take the above information into account when discussing future work in prioritizing the multipath work for QUIC", but the complete specification of ATSSS in the relevant 3GPP architecture documents reflects the complexity of 3GPP networks, the authors of this document worked to provide a high level overview of the parts of ATSSS that would be impacted by IETF

protocol design, in terminology that is more familiar to IETF participants.

1.1. Notes for Readers

We provide a high-level overview of ATSSS in [Section 5](#), describe the current ATSSS version in [Section 7](#), describe our thoughts about a QUIC-based version of ATSSS in [Section 7.1](#), and conclude with our understanding of the gaps between QUIC version 1 and what a QUIC-based version of ATSSS will require in [Section 8](#).

This document is an informational Internet-Draft from individuals, and does not carry any special status within the IETF.

This document abstracts considerable architectural detail that is available in 3GPP specifications. The goal is to make this overview more accessible for IETF readers who are not familiar with 3GPP 5G architecture.

This document makes references to Internet-Drafts that are not as mature as the core QUIC Internet-Drafts, and in some cases, have not been adopted as Working Group drafts yet, although all are within existing and proposed IETF working group charters. The goal is to give 3GPP readers the most up-to-date understanding of what is possible.

2. Introduction to Access Traffic Steering, Switching, and Splitting (ATSSS)

Mobile devices such as laptops, smartphones, tablets support multiple network interfaces that may attach to different networks. Over the years, various techniques have been proposed to support such multi-interfaced devices (e.g., Shim6 [[RFC5533](#)], Mobile IPv6 [[RFC6275](#)], Proxy Mobile IPv6 [[RFC5213](#)], or Multipath TCP [[RFC6824](#)]).

Users of these devices have different expectations concerning the utilization of available network connectivity (and thus their different network interfaces). For simplicity, we consider a smartphone that is equipped with a Wireless LAN (WLAN) interface and a cellular interface, but the discussion below can be generalized to any device with multiple network interfaces that support IP.

Some users of these smartphones want to offload most or all of their traffic onto the WLAN when the WLAN is available while expecting seamless handovers when it is not available. For example, when they move out of the reach of their home WLAN, they expect that the established flows (e.g., TCP connections, UDP flows) will continue over the cellular interface without any interruption (called "session

continuity" in 3GPP). As the devices are assigned different IP addresses over WLAN and the cellular networks, this seamless handover requires some specific assistance from the network. The current utilization of Multipath TCP on Apple smartphones is an example of this use case [[IETFJ16](#)].

Other users want to load balance their flows over the different available networks, e.g., by sending a delay-sensitive flow over cellular and a long download over the WLAN network. Several smartphones enable applications to indicate their preferences when using available networks. This steering policy can be managed by the smartphone, but flows need to continue after a handover.

Still other users may want to combine the resources provided by the cellular and the WLAN networks to improve the up and download throughput performance of individual flows. The GiGA LTE and GiGA 5G services deployed using Multipath TCP in South Korea are examples of this use case [[IETFJ16](#)].

To support these different use cases in 5G networks, 3GPP is defining the Access Traffic Steering, Switching and Splitting (ATSSS) service [[TS23501](#)]. This work is further adopted by the Broadband Forum to provide similar capabilities to residential gateways equipped with multiple access interfaces, in the continuity of the Hybrid Access Networks [[TR-348](#)].

In this document, we abstract many of the technical details of future 5G networks to explain the capabilities ATSSS needs, which may impact decisions about future work on IETF protocols.

3. Contribution and Discussion Venues for this draft.

(Note to RFC Editor - if this document ever reaches you, please remove this section)

This document is under development in the Github repository at <https://github.com/obonaventure/draft-quic-atsss-reqs>. Readers are invited to open issues and send pull requests with contributed text for this document.

Substantial discussion of this document should take place on the QUIC working group mailing list (quic@ietf.org). Subscription and archive details are at <https://www.ietf.org/mailman/listinfo/quic>.

4. Conventions, Terminology, and Definitions

This document makes use of 3GPP specific terms defined in [[RFC6459](#)], mainly the following ones:

- o Packet Data Network (PDN): is a packet-based network that either belongs to the operator or is external (such as the Internet or a corporate intranet). The user eventually accesses services in one or more PDNs. The operator's packet core networks are separated from packet data networks by User Plane Functions (UPFs).
- o UE (User Equipment): refers to the devices that are hosts with the ability to obtain Internet connectivity via a 3GPP network.
- o User Plane: refers to data traffic and the required sessions for the data traffic. In practice, IP is the only data traffic protocol used in the user plane.

Also, the document uses the following additional terms:

- o Protocol Data Unit (PDU) Session: An association between the UE and the Data Network (DN) to carry the user data/traffic.
- o PDU Connectivity Service: A service that provides exchange of PDUs between an UE and a Data Network.
- o Multi-access PDU (MA-PDU) Session: A PDU session that has simultaneously user plane resources assigned on 3GPP and non-3GPP access networks.
- o User Plane Function (UPF): A logical function in the 5G core network that provides the interconnect point between the mobile infrastructure and the Data Network (DN) and anchor point for Protocol Data Unit (PDU) Sessions to enable mobility.
- o Data Network Name (DNN): is a Fully Qualified Domain Name (FQDN) and resolves to a set of gateways in an operator's network. DNN is used for the selection of the UPF(s) for a PDU Session.
- o 5G Core (5GC) network: Refers to the part of the 5G System which is independent of the access technology used by an UE (e.g., cellular, WLAN) [[TS23501](#)]. A 5G Core network can be reached via one or more access networks.
- o 3GPP access network: Refers to a radio access network used by an UE to reach a 5G Core network. In such case, the UE uses an access technology that is specified by 3GPP.

- o Non-3GPP access network: Refers to an access network (e.g., WLAN) that is not a 3GPP access network and which is used by an UE to connect to a 5G Core network.
- o 5G control plane: Denotes the 5G control management component of use plane resources (e.g., forwarding policies).

5. High Level ATSSS Overview

The 5G Core supports a service that provides exchange of data between a User Equipment (UE) and a data network (referred to as Packet Data Network (PDN)) identified by a Data Network Name (DNN). This connectivity service, called the Protocol Data Unit (PDU) Connectivity Service, is realized via 'PDU Sessions' that are established upon request from User Equipment (UE) when the UE first connects to that network. The type of PDU Session can be IPv4, IPv6, IPv4v6, Ethernet or Unstructured.

It is out of the scope of this document to provide a comprehensive overview of 5G System (5GS) architecture. In particular, this document does not describe how PDU Sessions are established, and thus how IP addresses/prefixes are assigned to requesting UEs.

An UE can be provided a multi-access PDU Connectivity Service. That is, an UE can exchange data with a PDN by using a "3GPP access network" and a "non-3GPP access network" (often a WLAN). This is realized using the ATSSS service that is provided in the 5G core network control plane and user plane. The user plane part of the ATSSS functionality is contained in the User Plane Function (UPF) that manages the UE's PDU session.

5.1. Reference Architecture

To understand the operation of the ATSSS service, it is useful to consider the reference environment shown in Figure 1. An UE is attached to two different access networks (Access Net A and Access Net B). Each of these two networks is potentially shared with other users, so the bandwidth available from each network varies over time. These fluctuations in bandwidth are managed by using congestion control schemes.

One of these access networks is managed by a 5G provider according to the 3GPP specifications. The second network is potentially managed by a different organization. It is important to note that in this second case, there is an IPsec tunnel between the UE and a dedicated device in the 5G network (not shown in Figure 1). A dedicated IP address is assigned by means of Internet Key Exchange version 2 (IKEv2) to the UE to access the 5G Core via this second network.

The UE interacts with a distant server through a User Plane Function hosting the ATSSS functionality. This UPF is called hereafter ATSSS UPF. The ATSSS UPF enables the UE to use a transport protocol that supports the different access networks even if the server does not support it (i.e., does not use a multipath-capable transport protocol).

The 5G control plane provides the UE and the ATSSS UPF with rules as discussed in [Section 5.4](#). Note that, as per 3GPP definition, the rules provided to the UE are called ATSSS Rules, while the ATSSS information provided to the UPF is carried via Multi-Access Rules, but for simplicity this document uses the term "ATSSS rules" for the ATSSS information provided to both UE and UPF.

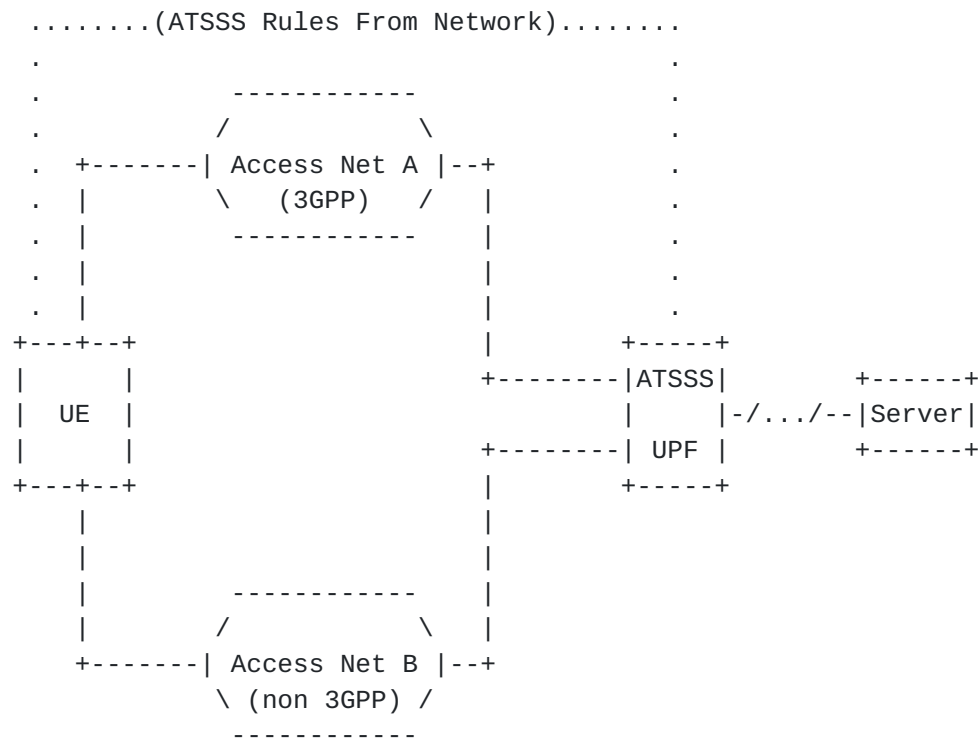


Figure 1: Simplified Reference Architecture for ATSSS

5.2. External IP Addresses Used by the ATSSS UPF

Depending on the ATSSS mode, the ATSSS UPF may translate the access network specific addresses into an address called the Multi-Access (MA) IP address and vice versa. Such an address is also directly assigned to the UE (that is, packets that are not eligible for the ATSSS service are sourced by the UE with that IP address).

Absent any coordination between the UE and the ATSSS UPF (e.g., assignment of port ranges), the same IP address and port number could

be used simultaneously by the ATSSS UPF and the UE; which is problematic.

To avoid such issue, the ATSSS UPF may be configured with a pool of IP addresses that it can use in the Internet-facing interfaces instead of preserving the MA IP address assigned to the UE. How that pool is used is deployment- and implementation-specific.

5.3. ATSSS Modes

3GPP defines the following procedures [[TS23501](#)] that are applicable between "3GPP access" and "non-3GPP access" networks:

- o Access Traffic Steering: selection of an access network for a new data flow and the transfer of the traffic of that data flow over the selected access network.
- o Access Traffic Switching: migration of all packets of an ongoing data flow from one access network to another access network. Only one access network is in use at a time, but this still ensures session continuity.
- o Access Traffic Splitting: forwarding the packets of a data flow across multiple access networks simultaneously.

Techniques to provide ATSSS are classified by the 3GPP into two flavors: (1) higher-layer techniques which operate above the IP layer (e.g., MPTCP), and (2) lower-layer techniques which operate below the IP layer.

5.4. ATSSS Rules

The 5G control plane provides the UE and the ATSSS UPF with rules that specify which flows are eligible to the ATSSS service (i.e., by mapping them to a Multi-Access PDU Session). Once a Multi-Access PDU Session has been established, a set of rules are then delivered to both the UE and the ATSSS UPF in order to enable consistent treatment of the flows by both the UE and the ATSSS UPF within the Session.

The traffic that matches an ATSSS rule can be distributed among the available access networks following one of these modes:

- o "Active-Standby": The traffic associated with the matching flow will be forwarded via a specific access (called 'active access') and switched to another access (called 'standby access') when the active access is unavailable.

- o "Smallest Delay": The traffic associated with the matching flow will be forwarded via the access that presents the smallest RTT. To that aim, specific measurements are conducted by the UE and a dedicated function co-located with the ATSSS UPF.
- o "Load-Balancing": The traffic associated with the matching flow will be distributed among the available access networks following a distribution ratio (e.g., 30% via a first access, 70% via a second access).
- o "Priority-based": For this mode, accesses are assigned priority levels that indicate which access to be used first. Concretely, the traffic associated with the matching flow will be steered on the access with a high priority till congestion is detected, then the overflow will be forwarded over a low priority access.

In order to provide the above-mentioned steering modes, measurement information about the current network state of each path is needed. Often if a multipath capable protocol is used, the measurements are available as part of the protocol itself. For ATSSS approaches where this is not the case, a dedicated protocol called Performance Measurement Function (PMF) protocol can be used. This protocol is enabled between the UE and the UPF in order to provide RTT measurements and report access availability/unavailability by the UE to the UPF.

These modes of operations can be met by using multipath protocols such as MPTCP [[RFC6824](#)] to select the different paths between the UE and the ATSSS. Such protocols usually include two types of mechanisms to control the utilization of the paths: (i) a path manager and (ii) a packet scheduler [[RFC8041](#)]. The path manager decides when a new subflow needs to be established over a path while the packet scheduler selects the subflow over which the next packet will be sent. A more detailed description of several packet schedulers may be found in [[I-D.bonaventure-iccr-g-schedulers](#)].

The "Active-Standby" mode can be implemented using a path manager that tries to use the active access and switches to the standby one after a certain number of retransmissions. This mode of operation is similar to the utilization of Multipath TCP on iOS smartphones [[RFC8041](#)].

The "Smallest Delay" mode can be implemented using a path manager that establishes a subflow over both paths and a packet scheduler that measures their RTTs and prefers the one having the lowest RTT. These path manager and scheduler are similar to those used by Multipath TCP on Linux [[RFC8041](#)].

The "Load-Balancing" mode would use the same path manager with a weighted round-robin scheduler.

The "Priority-based" mode can be implemented using a path manager that is similar to the one used by the "Active-Standby" one, but which reacts faster and a packet scheduler that prefers the high priority path. These path manager and scheduler are similar to the ones used in deployed Hybrid Access networks [[Hybrid](#)].

6. ATSSS Phases

We first describe in [Section 7](#) ATSSS as specified in Release 16 (called hereafter, ATSSS Phase 1) that uses Multipath TCP [[RFC6824](#)] and the 0-RTT Convert Protocol [[I-D.ietf-tcpm-converters](#)] to handle TCP traffic. We then discuss in [Section 7.1](#) the data plane requirements for Phase 2 of the ATSSS specification that 3GPP plans for Release 17. More details about 3GPP releases can be found at: <https://www.3gpp.org/specifications/Releases>.

7. ATSSS Phase 1: Support for TCP

For ATSSS with Multipath TCP functionality, a client with two interfaces connected to two disjoint access networks (in this case, Access Net A and Access Net B) uses MPTCP to reach an MPTCP proxy over either, or both, of the access networks. This allows the client to communicate with a server which does not support MPTCP.

During the attachment of an ATSSS-capable UE to the network, the UE may retrieve the MPTCP proxy information: an IP address, a port number, and the type of proxy. In the current release, the mandatory MPTCP proxy type is the "Transport Converter" [[I-D.ietf-tcpm-converters](#)].

Also, both the MPTCP Client and MPTCP proxy are configured with ATSSS rules from the network that govern how the multiple network paths between the MPTCP Client and MPTCP proxy are used. This relationship is shown using "." between the MPTCP Client and MPTCP Proxy in Figure 2.

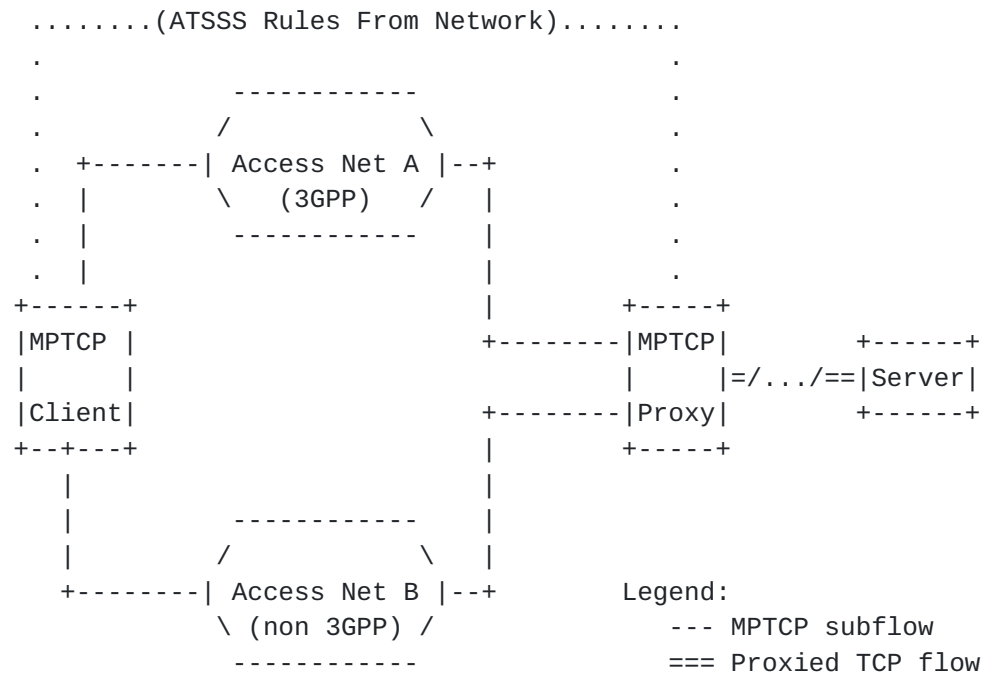



Figure 2: Simplified Reference Architecture for ATSSS with Multipath TCP

An ATSSS-capable UE can make use of the MPTCP functionality by establishing MPTCP-assisted connections via the MPTCP proxy relying upon the Convert Protocol [[I-D.ietf-tcpm-converters](#)]. The UE behaves as a "Client", while the MPTCP proxy behaves as a Transport Converter [[I-D.ietf-tcpm-converters](#)].

The UE then sends packets bound to connections matching an ATSSS rule to the provisioned Transport Converter and destination port number. Concretely, the UE initiates the MPTCP connection towards the Transport Converter and indicates the IP address and port number of the Server within the connection establishment packet (i.e., in the payload of the SYN sent to the Transport Converter). Doing so enables the Transport Converter to immediately initiate a connection towards that Server, without experiencing an extra delay. The Transport Converter waits until the receipt of the confirmation that the Server agrees to establish the connection before confirming it to the Client.

A flow example of an MPTCP-proxied connection is shown in Figure 3. This example assumes that the Server is not MPTCP-aware. The instructions included in the matching ATSSS rule will be followed for the management of the MPTCP connection (including the selection of the access network to establish the first subflow).



Legend:

[]: Convert Protocol TLVs

MPC: MP_CAPABLE option [[RFC6824](#)]

Figure 3: An Example of MPTCP-proxied Connection Matching an ATSSS Rule.

This approach provides 0-RTT (Zero Round-Trip Time) conversion service since no extra delay is induced by the Convert protocol compared to connections that are not proxied. Also, the Convert Protocol does not require any encapsulation (so, no tunnels). The UE and the MPTCP proxy track the performance of the access networks by leveraging MPTCP's internal mechanisms including congestion control and round-trip-time measurements. MPTCP uses this performance information to support splitting and switching.

If the server supports MPTCP, the Convert Protocol provides an option for clients to "opt-out". In such case, an MPTCP connection is directly established between the client and the server. Given that few servers are MPTCP capable, relying on the ATSSS service is the only option for UEs to make use of available multiple paths to most servers simultaneously.

7.1. ATSSS Phase 2: Adding Non-TCP Support

The MPTCP-based ATSSS approach discussed in the previous section is specific to TCP, obviously. Therefore it does not support non-TCP traffic, such as UDP, QUIC [[QUIC-Deployment](#)] or IPsec and Datagram Transport Layer Security (DTLS) Virtual Private Network (VPN) services.

As the share of these protocols grows, mainly driven by QUIC deployment, a future ATSSS needs to extend beyond supporting TCP only. The seamless handover provided by ATSSS is particularly useful for real-time traffic (e.g., voice or video calls),

Several proposals to carry non-TCP traffic have been discussed, including using TCP [[I-D.boucadair-mptcp-plain-mode](#)] or defining

multipath extensions to DCCP [[I-D.amend-tsvwg-multipath-dccp](#)]. The work within 3GPP now focuses on using QUIC as the baseline for ATSSS Phase 2.

7.1.1.1. QUIC and Multihoming

For non-TCP traffic, QUIC is already the dominant part of UDP traffic. A solution as realized with the Convert Protocol for (MP)TCP is not possible for QUIC as a QUIC connection cannot be intercepted and converted as in the ATSSS architecture for (MP)TCP (Figure 2). For (MP)TCP only the transport protocol (TCP) is intercepted, while transport security provided by Transport Layer Security (TLS) on top of TCP stays in tact. For QUIC transport, security is integrated into the transport protocol and thus cannot be intercepted.

Some ATSSS modes can be natively supported by the base QUIC specification for QUIC flows. For example, the "Active-Standby" and "Smallest Delay" steering modes can be supported directly between an UE and a QUIC server without any assistance from the network other than the performance measurement information.

Further, QUIC provides a feature called connection migration (Section 9 of [[I-D.ietf-quic-transport](#)]) that makes it possible to move a QUIC connection from one path/IP address to another without terminating and reestablishing the connection. Connection migration can further enable traffic switching but does not support traffic splitting as only one path can be used simultaneously.

An extension to QUIC to support simultaneous use of multiple paths is proposed in [[I-D.deconinck-quic-multipath](#)]. However, similar to a native MPTCP connection, a (MP)QUIC connection initiated between the UE and a server without the ATSSS UPF assistance cannot benefit from any direct application of the ATSSS steering methods based on network input given that the steering policy as currently defined in ATSSS is local to the UE and the ATSSS UPF and there are no means to signal that policy to a remote server.

Network input can be especially beneficial for cases such as:

- o avoiding unnecessary use of user quota if one of the access networks is subject to volume-based quota.
- o avoiding frequent connection migration if both access networks could be used to forward packets (each with a distinct source IP address).

7.1.2. QUIC as an ATSSS Data Plane Protocol

This section elaborates how non-TCP traffic (UDP or a subset like QUIC) can be encapsulated into a tunnel between the UE and the ATSSS UPF to enable ATSSS for all traffic, not only TCP or QUIC. This section discusses to what extent QUIC can be used as a tunneling protocol for the ATSSS service and whether gaps are found.

When tunneling non-TCP (e.g., UDP, IP) over QUIC the Unreliable Datagram Extension [[I-D.ietf-quic-datagram](#)] can be used. Each data packet would be transported unreliably as a datagram over a QUIC connection. Transporting these datagrams unreliably as would be done in IPsec or DTLS-based tunnels is especially important for flows that do not require reliable delivery and would suffer from unnecessary delays caused by the retransmissions used to support reliability. QUIC datagrams are congestion-controlled, but since the latency between the UE and the ATSSS UPF is small compared to the end-to-end latency, having second, local, congestion control loops should not impact the end-to-end congestion control negatively.

This document discusses three approaches:

- o Use of QUIC version 1 with the Unreliable Datagram Extension [Section 7.1.3](#)
- o Use of QUIC version 1 with the Unreliable Datagram Extension but with one QUIC connection over each access network [Section 7.1.4](#)
- o Use of a single Multipath QUIC connection [[I-D.deconinck-quic-multipath](#)], with the Unreliable Datagram Extension, over all access networks [Section 7.1.5](#).

7.1.3. Single QUICv1 Tunnel with Unreliable Datagram Extension and Connection Migration

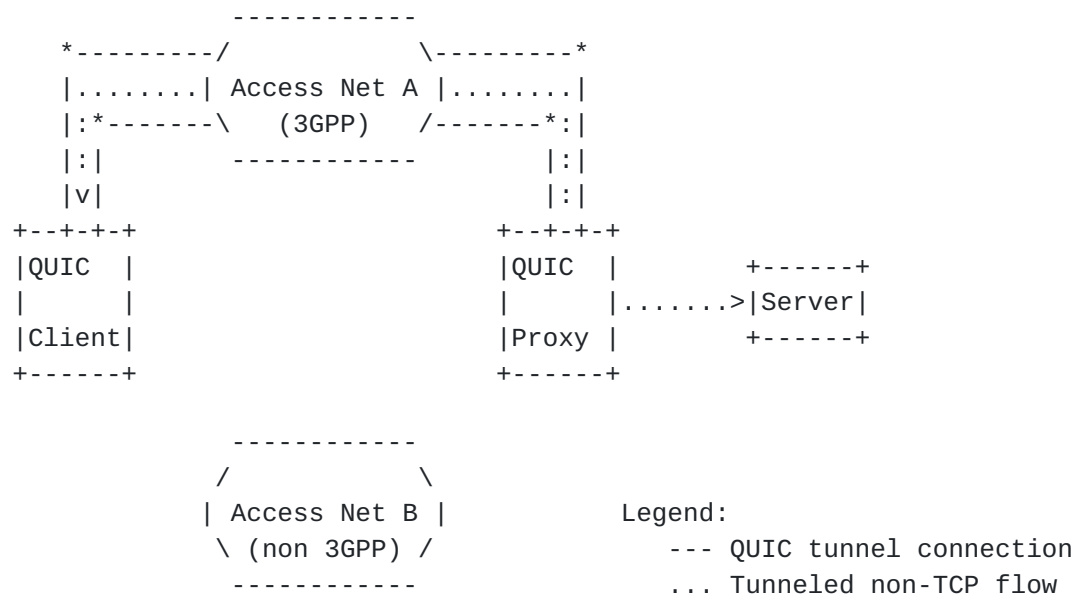


Figure 4: Single QUICv1 tunnel

QUIC can be used as a tunneling protocol between the UE and the ATSSS UPF. Use of the Unreliable Datagram Extension avoids unnecessary delays due to local retransmissions and, more important, subsequent head-of-line blocking.

In this case, there is (only) one QUIC connection between the UE and the ATSSS UPF for a given flow. And as such, that QUIC connection uses only one access network at a time. However, given the connection migration capability of QUIC, the QUIC connection could be moved to another access network, e.g., when the network indicates that the currently used access network would go away or that its capacity becomes limited. The UE can then initiate the connection migration to start the path validation process as specified in Section 9 of [I-D.ietf-quic-transport]. When, and how, to switch over depends on the rules provided by the network (Section 5.4) and the performance measurements that are accessible to both the UE and the ATSSS UPF. Note that connection migration can only be at the initiative of the UE as per Section 9 of [I-D.ietf-quic-transport]. This means that the UPF cannot make use of a second access network upon failure or degradation observed on a first access network.

It is thus possible to support the switching and steering functions of ATSSS, but splitting cannot be supported.

Path validation induces a delay when switching as packets are buffered. Further, congestion control state is reset on a new path and needs to ramp up. When frequent handovers happen, splitting traffic over multiple paths simultaneously can be beneficial for a

smooth user experience. Further, of course, the ability to use multiple paths simultaneously also increases the maximum capacity available, which can also be beneficial in some cases.

This option is not considered a valid approach for the ATSSS Phase 2 discussed within 3GPP.

7.1.4. Multiple QUICv1 Tunnels with Unreliable Datagram Extension and Connection Migration

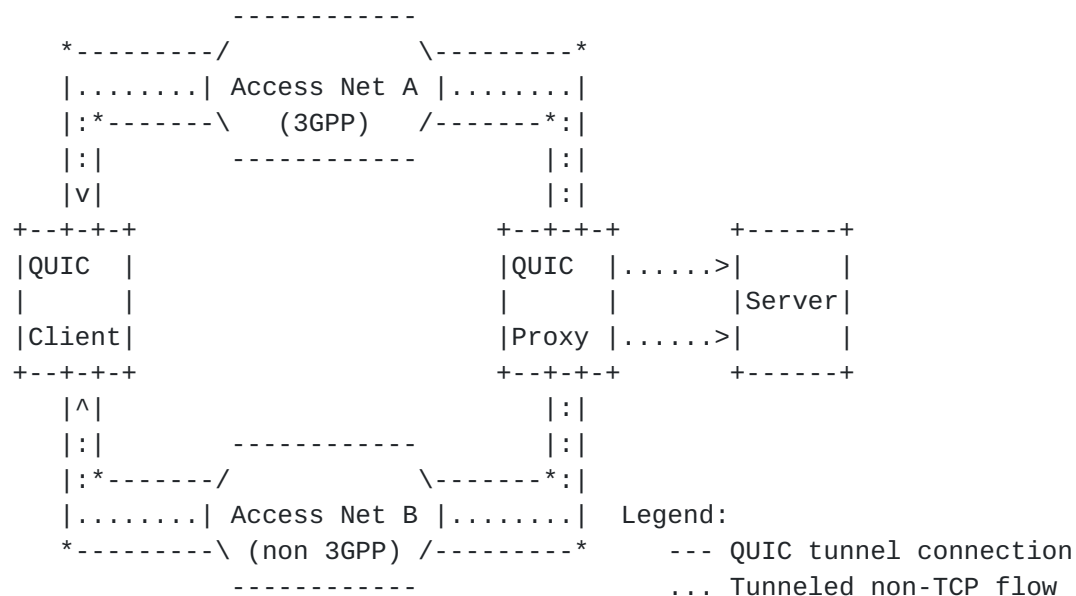


Figure 5: Traffic steering of two end-to-end flows using multiple QUICv1 tunnels

Another approach is to use one QUIC connection over each access network. In this case, there are multiple QUIC connections that are used as tunnels from the UE to the ATSSS UPF to transport traffic that belongs to one or multiple flows. The UE and ATSSS UPF need to select one QUIC connection to forward each application data packet, based on the rules provided by the network.

This approach can support steering (by selecting just one QUIC connection for each flow), switching (by moving flows from one QUIC connection to another) as well as splitting when both tunnels are used simultaneously for different packets of the same flow.

However, this approach for splitting is challenging since data from the same flow are sent over different QUIC connections, again using unreliable datagram to minimize head-of-line blocking. Because both these QUIC connections are completely independent of each other and the paths on the different access networks may have different

latency, this approach would likely result in reordering of packets that belong to a split flow. If reordering should be avoided, this would require additional signaling between the UE and the ATSSS UPF, e.g., adding sequence numbers, and adding a reordering buffer and logic to both the UE and ATSSS UPF.

Furthermore, since the bandwidth available for each QUIC connection varies as a function of the congestion experienced over each access network, data sent over a congested connection could delay the delivery of subsequent data over another connection.

Experience with MPTCP on smartphones shows that its integrated mechanisms (e.g., congestion control, round-trip-time estimation, packet scheduler) are well suited to support splitting and switching. Providing a similar service over independent QUIC connections would require a complex application that would need to track the congestion window, round-trip-time, and loss characteristics of the underlying QUIC connections as well as some specific application layer signalling to glue the various QUIC connections together.

7.1.5. MP-QUIC Tunneling

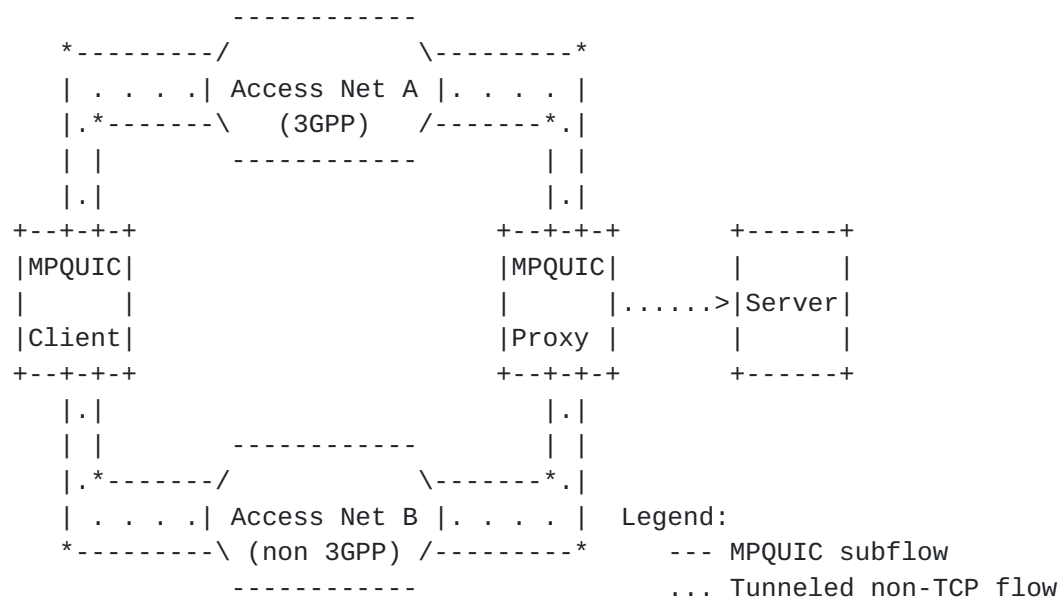


Figure 6: Traffic splitting using a Multipath QUIC tunnel

A third candidate solution is to leverage the ability of QUIC to support multiple streams and the Unreliable Datagram Extension, and to extend it with Multipath capabilities as described in [\[I-D.deconinck-quic-multipath\]](#).

In this case, there is a single (Multipath) QUIC connection between the UE and the ATSSS UPF. With a multipath transport, splitting is naturally supported.

Data sent over one access network can be retransmitted over the other if the first becomes congested. Some measurements and simulations have shown that Multipath QUIC provides similar performance as Multipath TCP when combining different access networks [[MPQUIC-Conext](#)].

Steering can be also supported, similarly to MPTCP. The path scheduler can map datagrams carrying an entire flow to one or another access network based on the provided rules.

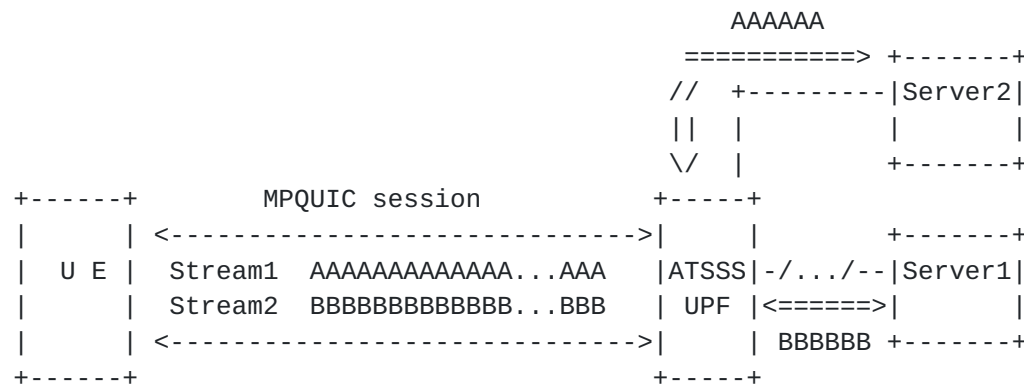
Switching can be improved by splitting traffic simultaneously over both links such that the congestion window of the new path can be open before the old path goes away. This makes handovers smoother. Experience with Multipath TCP on smartphones has shown that handovers are not a binary process. When a smartphone performs handovers, there are frequently short periods of time during which both networks are imperfect [[MPTester](#)]. Using use both networks simultaneously during these periods, improves user experience.

Furthermore, traffic can be better distributed among available paths based on available resources if one of the access networks fails or begins to become congested.

[7.2.](#) Mapping of Both TCP and Non-TCP to QUIC Streams and Datagrams

In ATSSS Phase 1, each TCP connection originated by the UE corresponds to one MPTCP connection that is terminated on the ATSSS UPF. With ATSSS Phase 2 using MPQUIC as a tunneling protocol, one can leverage the multi-stream capability of QUIC to carry the bytestreams of multiple TCP connections over a single MPQUIC session.

Focusing on the case where the UE maintains one QUIC connection with its ATSSS UPF, every TCP connection that it creates results in the creation of a new stream of the MPQUIC connection with the ATSSS UPF. Similarly the ATSSS UPF can map each incoming TCP connection from a remote host onto a stream of the MPQUIC connection. This is illustrated in Figure 7. Stream mapping is most beneficial for TCP, as it avoids reordering and TCP anyway requires in-order delivery. It is more appropriate to use QUIC with the Unreliable Datagram Extension for all other traffic.



<----> MPQUIC session

<====> TCP connection

Figure 7: ATSSS Phase 2 Maps Each TCP Connection on a Stream of the MPQUIC Connection Between the UE and the ATSSS UPF

Two application layer protocols are proposed to manage the mapping of TCP connections to QUIC streams as well as the transmission of UDP datagrams over QUIC datagrams:

(1) The proposed masque working group (wg) extends the HTTP/3 CONNECT method with a UDPCONNECT method to use QUIC as a tunnel for UDP traffic [[I-D.schinazi-masque-connect-udp](#)]. Another use case in scope for masque wg is IP proxying which can be used for tunneling and forwarding of TCP connections [[I-D.schinazi-masque-protocol](#)].

(2) QUIC Tunnel is a close proposal providing similar functionalities to MASQUE based on a binary protocol [[I-D.pirau-quit-tunnel](#)][I-D.pirau-quit-tunnel-tcp], and focusing on the ATSSS use case.

Both proposals rely upon single QUIC connections and inherit thus the same issues discussed in [Section 7.1.3](#) and [Section 7.1.4](#).

7.3. Encapsulation Overhead

In 3GPP architectures, a variety of encapsulations are already used to carry user data. The use of QUIC as a tunneling method for ATSSS will add some additional overhead. When the user data is forwarded over a non-3GPP network access, this overhead comes further in addition to an IPsec tunnel between the UE to the 5G Core Network which is already at least 142 octets for an IPv6 packet forwarded over a non-3GPP network access. In contrast, the MPTCP based mechanism in ATSSS Phase 1 does only add a minor overhead during

connection establishment, but no additional overhead during the rest of the connection.

More investigation is required to assess whether the ATSSS Phase 2 overhead is an issue. Solutions to optimize that overhead could be considered, if needed.

7.4. Multiple Encryptions

The use of QUIC as a tunneling protocol in ATSSS will add an additional layer of encryption. As there is already encryptions on other layers (e.g., IP Encapsulating Security Payload (ESP)) this leads to multiple layers of encryption, which might be redundant.

Such multiple encryptions will have performance implications on the UE, in particular. Means to optimize the various redundant encryptions are for further investigation in 3GPP.

7.5. Congestion Control in Congestion Control and Coexistence

Many applications that are sending unreliable traffic use various congestion control algorithms to detect congestion and adjust their sending rate based on inferred end-to-end latency and loss characteristics. Examples include Adaptive Video Streaming using e.g. SCREAM [[RFC8298](#)], or other media streaming applications that use QUIC as transport layer. When using QUIC version 1 with Unreliable Datagram Extension or Multipath QUIC as ATSSS solution, this leads to nested congestion control, where both inner and outer congestion control coexist. When using Multiple QUICv1 tunnels with Unreliable Datagram Extension or MP-QUIC tunneling, the packet scheduler could have an additional impact on the perceived end-to-end latency and loss due to the potential difference of the individual path characteristics.

When deploying ATSSS Phase 1 and Phase 2 in parallel, with the UE serving both reliable and unreliable flows, different congestion control algorithms may coexist on individual paths, which may lead to fairness issues or even meltdown effects.

We recognize that nested congestion control mechanisms (such as QUIC encapsulated in QUIC or SCREAM encapsulated in QUIC) as well as the coexistence of ATSSS Phase 1 and Phase 2 have implications that require further study.

7.6. Packet Order Reconstruction for (MP-)QUIC Splitting Mode

Both tunnel solutions in [Section 7.1.4](#) and [Section 7.1.5](#) allow the splitting mode for simultaneously sending data over multiple paths. In this case packet reordering can occur when a QUIC tunnel communication is split across paths with very different latencies. Generally, applications have to deal with packet reordering, since the best effort for Internet traffic has no guarantee to prevent it. However, in practice, packet reordering in the network is assumed to be very limited. Applications that require in-order delivery and e.g. rely on TCP or implement a similar mechanism itself can be sensitive to high amounts of reordering and experience decreased performance.

As such it is desirable that the respective receiver side of the tunnel termination points has the ability to reconstruct the original packet ordering. While in the case when using the MPTCP converter, losses are retransmitted quickly on the local segment, when the Unreliable Datagram Extension for QUIC is used, the reconstruction mechanism has to further account for packet loss which may occur for any of the paths within the QUIC tunnel. This can cause delays in the reordering logic, which in turn can have a negative impact on applications that do not require in-order delivery, such as real-time transmissions.

It is assumed that this is a solvable task similar to [\[MPDCCP-paper\]](#), but is probably left to the implementer, to take care. Even if the reconstruction of the packet order does not become a standardized part of the MP-QUIC in [Section 7.1.5](#), it possibly requires path sequencing and end-to-end sequencing.

8. QUICv1 Gap Analysis for ATSSS Phase 2

This section summarizes QUIC protocol capabilities that would be beneficial for ATSSS Phase 2, as described in [Section 7.1](#).

- o ATSSS Phase 2 is focused on transport services for Ethernet frames and IP packets (with the intention of supporting TCP, UDP, and UDP-encapsulated transport protocols such as QUIC). The discussed approaches are based on tunnelling Ethernet or IP directly over QUIC. The masque working group that is currently in the chartering validation process is scoped to cover UDP and IP proxying. While UDP proxying does cover the most important use case to support ATSSS for UDP/QUIC, a more generic solution based on IP proxying would simplify the ATSSS design. However, IP proxying is only considered at a later stage by the masque working group. Also, multipath mechanisms of QUIC are not covered in the proposed charter.

- o We envision the ability to select paths (steering), detect path failures and reroute traffic (switching), and forward packets on multiple active paths simultaneously (splitting), based on external policies, including active-standby, smallest delay, weighted load-balancing, and path selection based on assigned priorities, for the full range of encapsulated protocols in ATSSS Phase 2, similar to the abilities provided by Multipath TCP for ATSSS Phase 1. Splitting cannot be supported easily in the discussed QUICv1-based approaches. Multipath transport capability similar to Multipath TCP, as used in ATSSS Phase 1, would support splitting well. The QUIC working group is originally chartered to produce a multipath extension document by December 2021. Proposals exist, however, this work has been postponed to QUICv2 and discussion is still on-going if support will be kept in the charter.
- o While the base protocol for QUICv1 does not provide support for unreliable datagrams, an QUIC extension for datagram support has been adopted by the group and the QUIC working group is chartered to produce this capability by March 2021. This can be used to support for the additional user-plane protocols as envisioned in ATSSS Phase 2.
- o When QUIC is used as a tunneling protocol, nested congestion control mechanisms (such as QUIC encapsulated in QUIC) have implications that require further study.
- o When QUIC is used as a tunneling protocol, the complete ATSSS Phase 2 protocol stack would include encrypted headers at multiple layers. It needs further investigation if this is a problem for ATSSS, however, it is likely a problem that can be solved by the 3GPP. Likewise, the implications of the various encapsulation overhead is to be further assessed within 3GPP.

9. IANA Considerations

This document does not make any request to IANA.

10. Security Considerations

[Section 9 of \[RFC6459\]](#) provides an overview of security considerations in 3GPP networks. ATSSS Phase 1 data plane security considerations are documented in Section 9 of [\[I-D.ietf-tcpm-converters\]](#).

This document discusses the use of QUIC (including Multipath QUIC) as an additional ATSSS steering method. QUIC-specific security

considerations are discussed in Section 21 of [\[I-D.ietf-quic-transport\]](#).

This document does not specify specific mechanisms to use QUIC as a tunneling protocol towards an ATSSS proxy, as the intention of this document is to provide an informational overview of the ongoing work in 3GPP on ATSSS to support non-TCP, rather than discussing a detailed solution. Nevertheless, this document cites candidate solutions to provide such tunneling service. Security considerations specific to these solutions are provided below.

Multipath QUIC-specific security considerations can be found in Section 8 of [\[I-D.deconinck-quic-multipath\]](#).

[Section 6](#) of {I-D.ietf-quic-datagram} discusses security considerations specific to the use of the Unreliable Datagram Extension to QUIC.

[11.](#) Acknowledgements

Many thanks to Anna Brunstrom, Dieter Gludovacz, Dirk von-Hugo, Tim Costello, Stephen Johnson, and Florin Baboescu for the review, comments, and suggestions.

[12.](#) Document History

(Note to RFC Editor - if this document ever reaches you, please remove this section)

Version -00: initial submission

[13.](#) References

[13.1.](#) Normative References

[TS23501] 3GPP (3rd Generation Partnership Project), ., "Technical Specification Group Services and System Aspects; System Architecture for the 5G System; Stage 2 (Release 16)", 2019, <https://www.3gpp.org/ftp/Specs/archive/23_series/23.501/>.

[13.2.](#) Informative References

[atsssliaison]
3GPP (3rd Generation Partnership Project), ., "LS on need for Multi-Path QUIC for ATSSS", April 2020, <<https://datatracker.ietf.org/liaison/1678/>>.

- [Hybrid] Keukeleire, N., Hesmans, B., and O. Bonaventure, "Increasing broadband reach with Hybrid Access Networks", IEEE Communications and Standards Magazine, Vol. 4, Issue 1, March 2020, <<https://doi.org/10.1109/MCOMSTD.001.1900036>>.
- [I-D.amend-tsvwg-multipath-dccp]
Amend, M., Bogenfeld, E., Brunstrom, A., Kassler, A., and V. Rakocevic, "DCCP Extensions for Multipath Operation with Multiple Addresses", [draft-amend-tsvwg-multipath-dccp-03](#) (work in progress), November 2019.
- [I-D.bonaventure-iccrs-schedulers]
Bonaventure, O., Piraux, M., Coninck, Q., Baerts, M., Paasch, C., and M. Amend, "Multipath schedulers", [draft-bonaventure-iccrs-schedulers-00](#) (work in progress), March 2020.
- [I-D.boucadair-mptcp-plain-mode]
Boucadair, M., Jacquenet, C., Bonaventure, O., Behaghel, D., stefano.secci@lip6.fr, s., Henderickx, W., Skog, R., Vinapamula, S., Seo, S., Cloetens, W., Meyer, U., Contreras, L., and B. Peirens, "Extensions for Network-Assisted MPTCP Deployment Models", [draft-boucadair-mptcp-plain-mode-10](#) (work in progress), March 2017.
- [I-D.deconinck-quic-multipath]
Coninck, Q. and O. Bonaventure, "Multipath Extensions for QUIC (MP-QUIC)", [draft-deconinck-quic-multipath-04](#) (work in progress), March 2020.
- [I-D.ietf-quic-datagram]
Pauly, T., Kinnear, E., and D. Schinazi, "An Unreliable Datagram Extension to QUIC", [draft-ietf-quic-datagram-00](#) (work in progress), February 2020.
- [I-D.ietf-quic-transport]
Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport", [draft-ietf-quic-transport-28](#) (work in progress), May 2020.
- [I-D.ietf-tcpm-converters]
Bonaventure, O., Boucadair, M., Gundavelli, S., Seo, S., and B. Hesmans, "0-RTT TCP Convert Protocol", [draft-ietf-tcpm-converters-19](#) (work in progress), March 2020.

[I-D.piraux-quic-tunnel]

Piraux, M. and O. Bonaventure, "Tunneling Internet protocols inside QUIC", [draft-piraux-quic-tunnel-01](#) (work in progress), March 2020.

[I-D.piraux-quic-tunnel-tcp]

Piraux, M. and O. Bonaventure, "Tunneling TCP inside QUIC", [draft-piraux-quic-tunnel-tcp-00](#) (work in progress), March 2020.

[I-D.schinazi-masque-connect-udp]

Schinazi, D., "The CONNECT-UDP HTTP Method", [draft-schinazi-masque-connect-udp-00](#) (work in progress), April 2020.

[I-D.schinazi-masque-protocol]

Schinazi, D., "The MASQUE Protocol", [draft-schinazi-masque-protocol-01](#) (work in progress), March 2020.

[IETFJ16] Bonaventure, O. and S. Seo, "Multipath TCP Deployment", IETF Journal, Fall 2016 , 2016, <<https://www.ietfjournal.org/multipath-tcp-deployments/>>.

[MPDCCP-paper]

Amend, M., Bogenfeld, E., Cvjetkovic, M., Rakocevic, V., Pieska, M., Kassler, A., and A. Brunstrom, "A Framework for Multiaccess Support for Unreliable Traffic using Multipath DCCP", 2019 IEEE 44th Conference on Local Computer Networks (LCN) , 2019, <<https://doi.org/10.1109/LCN44214.2019.8990746>>.

[MPQUIC-Conext]

De Coninck, Q. and O. Bonaventure, "Multipath QUIC - Design and evaluation", Proceedings of the 13th international conference on emerging networking experiments and technologies 2017 Nov 28 (pp. 160-166). , 2017, <<https://multipath-quic.org/>>.

[MPTester]

De Coninck, Q. and O. Bonaventure, "MultipathTester - Comparing MPTCP and MPQUIC in Mobile Environments. In 2019 Network Traffic Measurement and Analysis Conference (TMA)", 2019, <<https://multipath-quic.org/multipathtester/2019/06/18/mnm-paper.html>>.

[QUIC-Deployment]

Kuehlewind, M., "Some updates on QUIC deployment numbers", 2019, <<https://datatracker.ietf.org/meeting/106/materials/slides-106-maprg-quic-deployment-update-00>>.

- [RFC5213] Gundavelli, S., Ed., Leung, K., Devarapalli, V., Chowdhury, K., and B. Patil, "Proxy Mobile IPv6", [RFC 5213](#), DOI 10.17487/RFC5213, August 2008, <<https://www.rfc-editor.org/info/rfc5213>>.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", [RFC 5533](#), DOI 10.17487/RFC5533, June 2009, <<https://www.rfc-editor.org/info/rfc5533>>.
- [RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", [RFC 6275](#), DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.
- [RFC6459] Korhonen, J., Ed., Soininen, J., Patil, B., Savolainen, T., Bajko, G., and K. Iisakkila, "IPv6 in 3rd Generation Partnership Project (3GPP) Evolved Packet System (EPS)", [RFC 6459](#), DOI 10.17487/RFC6459, January 2012, <<https://www.rfc-editor.org/info/rfc6459>>.
- [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>>.
- [RFC8041] Bonaventure, O., Paasch, C., and G. Detal, "Use Cases and Operational Experience with Multipath TCP", [RFC 8041](#), DOI 10.17487/RFC8041, January 2017, <<https://www.rfc-editor.org/info/rfc8041>>.
- [RFC8298] Johansson, I. and Z. Sarker, "Self-Clocked Rate Adaptation for Multimedia", [RFC 8298](#), DOI 10.17487/RFC8298, December 2017, <<https://www.rfc-editor.org/info/rfc8298>>.
- [TR-348] Broadband Forum, ., "Hybrid Access Broadband Network Architecture", 2016, <<https://www.broadband-forum.org/download/TR-348.pdf>>.

Authors' Addresses

Mohamed Boucadair (editor)
Orange
Clos Courtel
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Olivier Bonaventure (editor)
UCLouvain

Email: Olivier.Bonaventure@uclouvain.be

Maxime Piraux (editor)
UCLouvain

Email: Maxime.Piraux@uclouvain.be

Quentin De Coninck
UCLouvain

Email: quentin.deconinck@uclouvain.be

Spencer Dawkins (editor)
Tencent America

Email: spencerdawkins.ietf@gmail.com

Mirja Kuehlewind (editor)
Ericsson

Email: mirja.kuehlewind@ericsson.com

Markus Amend
Deutsche Telekom

Email: markus.amend@telekom.de

Andreas Kassler
Karlstad University

Email: andreas.kassler@kau.se

Qing An
Alibaba Group

Email: anqing.aq@alibaba-inc.com

Nicolas Keukeleire
Tessares

Email: nicolas.keukeleire@tessares.net

SungHoon Seo
Korea Telecom

Email: sh.seo@kt.com

