

Network Working Group
Internet-Draft
Intended status: Informational
Expires: June 10, 2017

M. Boucadair, Ed.
C. Jacquenet, Ed.
Orange
O. Bonaventure, Ed.
Tessares
W. Henderickx, Ed.
Nokia/Alcatel-Lucent
R. Skog, Ed.
Ericsson
December 7, 2016

Network-Assisted MPTCP: Use Cases, Deployment Scenarios and Operational Considerations

[draft-nam-mptcp-deployment-considerations-01](#)

Abstract

Network-Assisted MPTCP deployment models are designed to facilitate the adoption of MPTCP for the establishment of multi-path communications without making any assumption about the support of MPTCP by the communicating peers. MPTCP Conversion Points (MCPs) located in the network are responsible for establishing multi-path communications on behalf of endpoints, thereby taking advantage of MPTCP capabilities to achieve different goals that include (but are not limited to) optimization of resource usage (e.g., bandwidth aggregation), of resiliency (e.g., primary/backup communication paths), and traffic offload management.

This document describes Network-Assisted MPTCP uses cases, deployment scenarios, and operational considerations.

Requirements Language

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](#)].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Draft

Network-Assisted MPTCP

December 2016

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 June 10, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Terminology	4
3.	Network-Assisted MPTCP Use Cases	4
4.	Deployment Scenarios	5
4.1.	LTE/WLAN Aggregation	6
4.2.	Fixed/Wireless Access Aggregation	6
4.3.	Data Center	7
5.	Deployment & Operational Considerations	7
5.1.	MCP Location	8
5.2.	MCP Insertion in a Multipath Communication	8
5.2.1.	Explicit Mode (Off-path)	8
5.2.2.	Implicit Mode (On-path)	10
5.3.	Authorization	14
5.4.	MCP Behaviors	15
5.4.1.	Transparent MCP	15
5.4.2.	Non-Transparent MCP	17
5.5.	Address Family Considerations	19
5.6.	Policies & Configuration Parameters	19
5.6.1.	Towards End-to-End MPTCP Connections	19

5.6.2.	Traffic Distribution Scheme	22
5.6.3.	Flows Eligible to Multipath Service	22
5.6.4.	TCP Fragmentation	23
5.6.5.	DSCP Preservation	23
5.6.6.	Supported Transport Protocols	23

5.6.7.	Logging	23
6.	IANA Considerations	23
7.	Security Considerations	23
7.1.	Privacy	23
7.2.	Denial-of-Service (DoS)	24
7.3.	Illegitimate MCP	24
7.4.	High Rate Reassembly	24
8.	Contributors	24
9.	Acknowledgements	26
10.	References	26
10.1.	Normative References	26
10.2.	Informative References	26
	Authors' Addresses	28

[1.](#) Introduction

The overall quality of connectivity services can be enhanced by combining several access network links for various purposes – resource optimization, better resiliency, etc. Some transport protocols, such as Multipath TCP, can help achieve such better quality, but failed to be massively deployed so far.

The support of multipath transport capabilities by communicating hosts remains a privileged target design so that such hosts can directly use the available resources provided by a variety of access networks they can connect to. Nevertheless, network operators do not control end hosts while the support of MPTCP by content servers remains close to zero.

Network-Assisted MPTCP deployment models are designed to facilitate the adoption of MPTCP for the establishment of multi-path communications without making any assumption about the support of MPTCP capabilities by communicating peers. Network-Assisted MPTCP deployment models rely upon MPTCP Conversion Points (MCPs) that act on behalf of hosts so that they can take advantage of establishing communications over multiple paths.

Such MCPs can be deployed in CPEs (Customer Premises Equipment), as well in the network side. MCPs are responsible for establishing multi-path communications on behalf of endpoints.

This document describes Network-Assisted MPTCP uses cases ([Section 3](#)), deployment scenarios ([Section 4](#)), and operational considerations ([Section 5](#)).

[2.](#) Terminology

The reader should be familiar with the terminology defined in [\[RFC6824\]](#).

This document makes use of the following terms:

- o Client: an endhost that initiates transport flows forwarded along a single path. Such endhost is not assumed to support multipath transport capabilities.
- o Server: an endhost that communicates with a client. Such endhost is not assumed to support multipath transport capabilities.
- o Multipath Client: a Client that supports multipath transport capabilities.
- o Multipath Server: a Server that supports multipath transport capabilities. Both the client and the server can be single-homed or multi-homed. However, for the use cases discussed in this document, the number of interfaces on the endhosts is not relevant.
- o Transport flow: a sequence of packets that belong to a unidirectional transport flow and which share at least one common characteristic (e.g., the same destination address). TCP and SCTP flows are composed of packets that have the same source and destination addresses, the same protocol number and the same source and destination ports.

- o Multipath Conversion Point (MCP): a function that terminates a transport flow and relays all data received over it over another transport flow.

MCP is a function provided by the network operator that converts a multipath transport flow and relays it over a single path transport flow and vice versa.

3. Network-Assisted MPTCP Use Cases

The first use case is a Multipath Client that interacts with a Server. To benefit from the capabilities of its multipath transport protocol, the Multipath Client will interact with a Multipath Conversion Point (MCP) located in the network as illustrated in Figure 1.

```
C <=====>MCP <-----> S
+<=====>+
```

Legend:

<==>: MPTCP leg

Figure 1: A Multipath Client interacts with a Server through a Multipath Conversion Point

A similar approach consists of a Multipath Server that leverages its multipath capabilities when interacting with a Client as shown in Figure 2.

```
C <-----> MCP <=====> S
+<=====>+
```

Figure 2: A Client interacts with a Multipath Server through a Multipath Conversion Point

The third use case is when a Client interacts with a Server and the corresponding communication is deemed eligible to multi-path forwarding. In this case, two Multipath Conversion Points are used.

The upstream MCP converts the (single path) transport flow initiated by the Client into a multipath transport flow towards a downstream MCP (called, network-located MCP). This downstream MCP converts the multipath transport flow received from the upstream MCP in a single path transport flow forwarded to the destination Server. The end-to-end transport flow between the Client and the Server is thus decomposed into three flows as shown in Figure 3: A (single path) transport flow between the Client and the upstream MCP, a multipath transport flow between the upstream and the downstream MCPs, and a single path transport flow between the downstream MCP and the Server.




Figure 3: A Client interacts with a Server through two Multipath Conversion Points

4. Deployment Scenarios

This section discusses some deployment scenarios related to Network-Assisted MPTCP designs.

4.1. LTE/WLAN Aggregation

This deployment scenario is considered by mobile operators so that they can propose their customers to aggregate LTE resources with WLAN resources.

As depicted in Figure 4, the mobile terminal (UE, User Equipment) is MPTCP-capable. The MCP function is enabled in the network to terminate MPTCP connections (e.g., in the PDN Gateway, a dedicated service function located at the (S)Gi interface, co-located with a TCP proxy, etc.).

This deployment scenario is an implementation of the use case depicted in Figure 1.

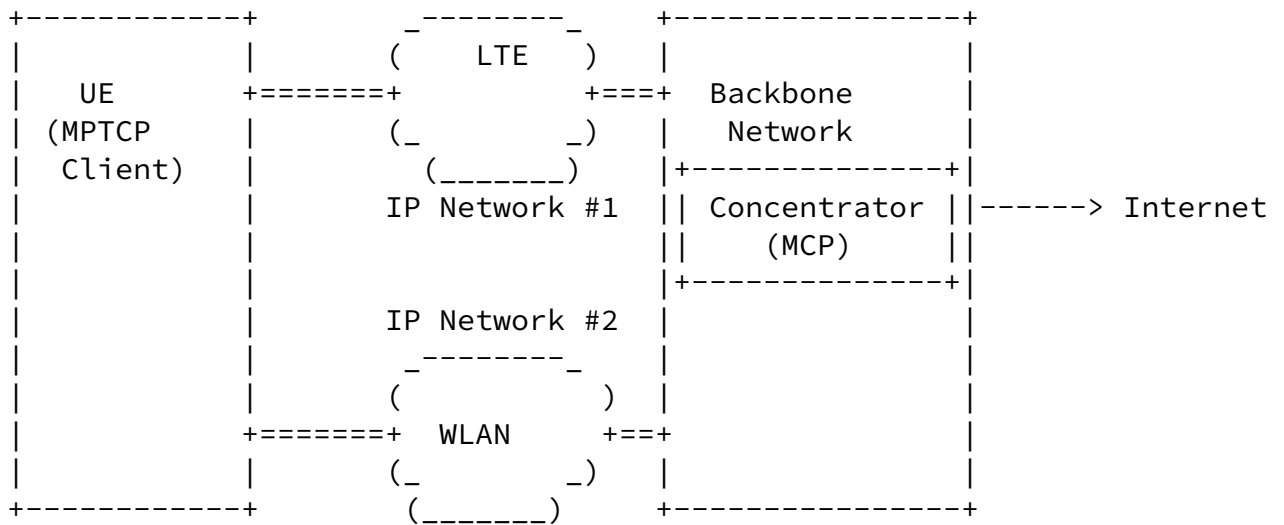
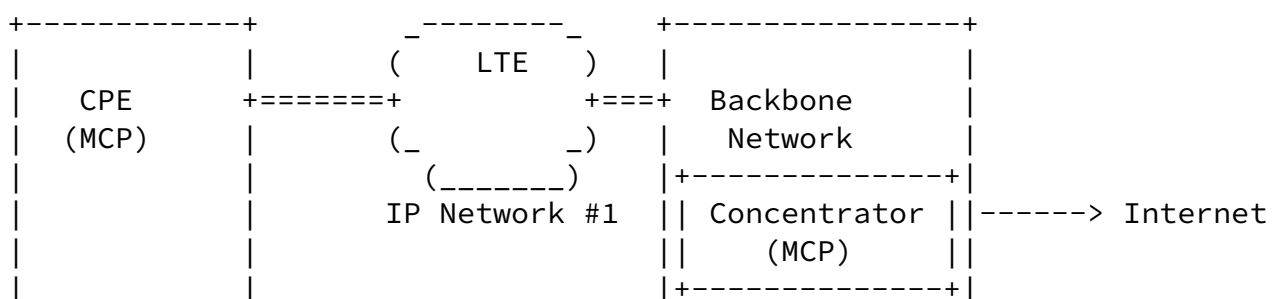


Figure 4: Network-Assisted MPTCP LTE/WLAN Aggregation (Host-based model)

4.2. Fixed/Wireless Access Aggregation

One of the promising deployment scenarios for Multipath TCP is to enable a CPE that is connected to multiple access networks (e.g., DSL, LTE, WLAN) to optimize the usage of such resources. This deployment scenario, called Hybrid Access, relies upon MCPs located in both the CPE and the network (Figure 5). The latter plays the role of an MPTCP concentrator. Such concentrator terminates the MPTCP sessions established from CPEs, before redirecting traffic into legacy TCP sessions.

This deployment scenario is an implementation of the use case depicted in Figure 2.






Figure 5: Network-Assisted MPTCP Fixed/Wireless Access Aggregation

For mobile operators that provide CPE-based mobile broadband services, LTE and WLAN resources can be aggregated by means of MPTCP. In such deployment scenario, the MCP function is enabled in the CPE and in the network.

4.3. Data Center

As discussed in Section 2.1 of [[I-D.ietf-mptcp-experience](#)], MPTCP is being contemplated for deployment within data centers. Such deployments may adhere to the use cases defined in Figure 2 or Figure 3. One or two MCPs may be enabled at the data center segment. The objective is to optimize the resources usage within the data center infrastructure.

The presence of an MCP is transparent to the client side, nevertheless the origin client's IP address may be hidden to the ultimate server. Enforcing per-host policies at the server's side may be problematic if all connections are bound to an MCP's IP address. To mitigate this problem, MCPs may insert the TCP Option for Host Identification [[RFC7974](#)].

5. Deployment & Operational Considerations

5.1. MCP Location

The location of MCPs is deployment-specific. Network Providers may choose to adopt centralized or distributed designs. Nevertheless, in order to take advantage of MPTCP, the location of an MCP should not jeopardize packet forwarding performance overall.

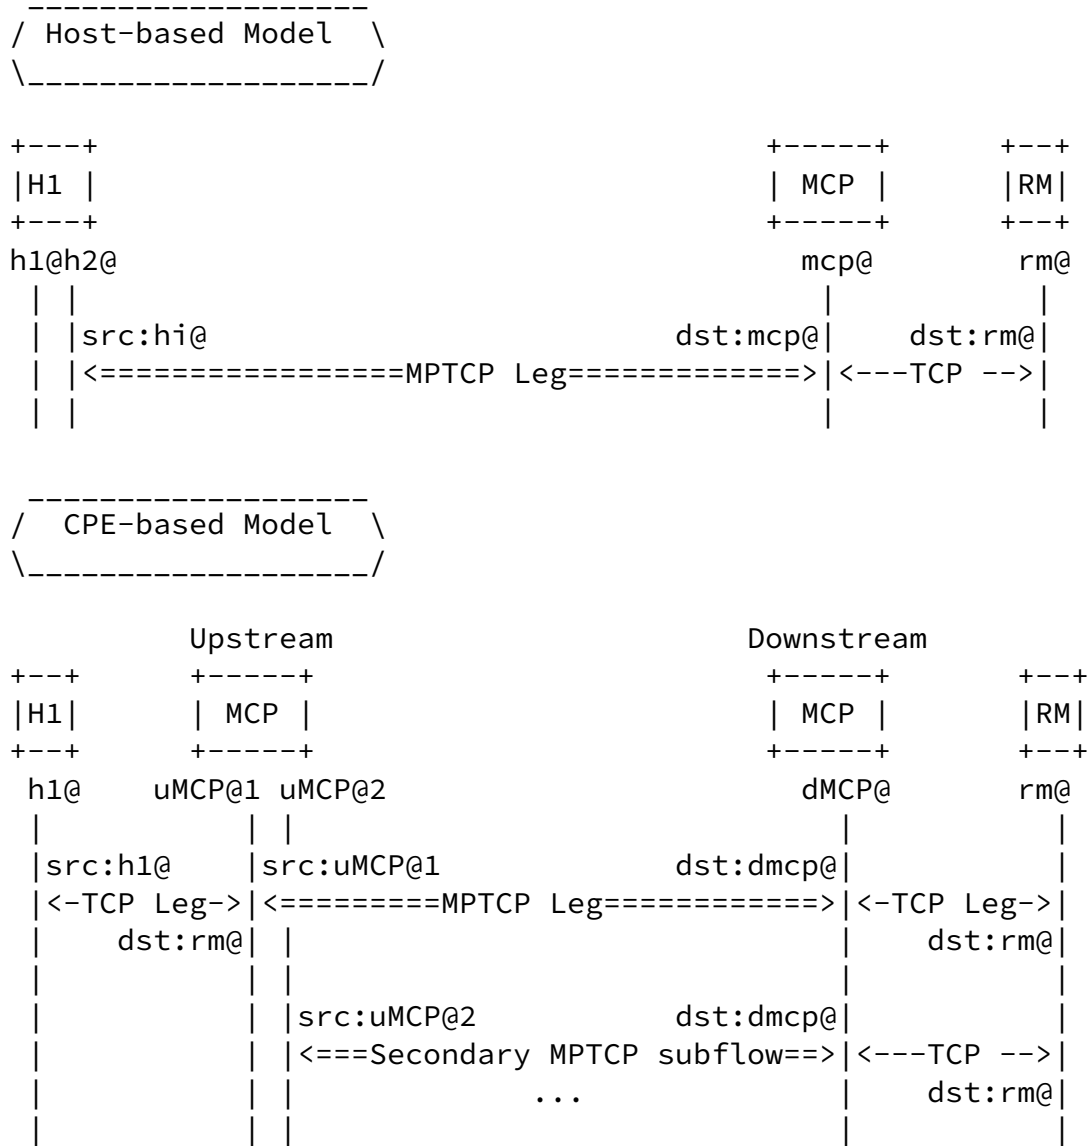
[5.2.](#) MCP Insertion in a Multipath Communication

Two deployment scenarios can be considered for involving an MCP in the communication path. These scenarios are described below.

[5.2.1.](#) Explicit Mode (Off-path)

This scenario assumes that the IP reachability information of an MCP is explicitly configured on a device, e.g., by means of a specific DHCP option [[I-D.boucadair-mptcp-dhc](#)]. A device can be a CPE or a host.

MPTCP connections are established explicitly using the address(es) of the MCP (Figure 6). In order to forward packets to their ultimate destination, the MCP is provided during the connection establishment with the destination IP address (and optionally destination port numbers). Typically, this is achieved thanks to the use of the MP_CONVERT Information Element (IE) defined in [[I-D.boucadair-mptcp-plain-mode](#)]. Figure 6 illustrates the flow exchange to establish a communication with a legacy server (RM).



Legend:

RM: A remote machine.

Figure 6: Sample Connection Establishment (Explicit Mode)

This scenario aims to avoid any adherence of the Network-Assisted MPTCP procedure and the underlying routing and forwarding policies. Furthermore, this scenario allows for more flexibility in terms of mounting MPTCP subflows as it does not require any specific order in the establishment of subflows among available interfaces.

Because the MCP's reachability information is explicitly configured on the device, means to guarantee successful inbound MPTCP connections can be enabled in the device to instruct the MCP to

maintain active bindings so that incoming packets can be successfully redirected towards the appropriate device.

[5.2.2.](#) Implicit Mode (On-path)

Unlike the explicit mode, the implicit mode assumes that the MCP is located on a default forwarding path (primary path). As such, the first subflow must always be placed over that primary path so that the MCP can intercept MPTCP flows. Once intercepted, the MCP advertises its reachability information by means of MPTCP signals (MP_JOIN or ADD_ADDR). Figure 7 illustrates the flow exchange to establish a communication with a legacy server (RM).

Note that a standard MPTCP implementation will try to establish other subflows using `rm@`. In order to prevent such behavior, H1 is instructed by some means to prevent establishing additional subflows to `rm@`. For example, the MCP may set the C-bit in MP_CAPABLE option ([\[I-D.ietf-mptcp-rfc6824bis\]](#)) it returns to H1.

Internet-Draft

Network-Assisted MPTCP

December 2016

```

-----
/ Host-based Model \
-----

```

+-----+	+-----+	+---+
H1	MCP	RM
+-----+	+-----+	+---+
h1@h2@	mcp@	rm@
src:h1@	dst:rm@	dst:rm@
<=====Initial MPTCP subflow=====> <---TCP -->		
	...	
src:h2@	dst:mcp@	dst:rm@
<=====Secondary MPTCP subflow=====> <---TCP -->		
	...	

```

-----
/ CPE-based Model \
-----

```

Upstream			Downstream		
+---+	+-----+		+-----+	+---+	
H1	MCP		MCP	RM	
+---+	+-----+		+-----+	+---+	
h1@	uMCP@1 uMCP@2		dMCP@	rm@	
src:h1@	src:uMCP@1		dst:rm@	src:uMCP@1	
<-TCP Leg-> <=====MPTCP Leg=====> <-TCP Leg->				dst:rm@	
	dst:rm@				
	src:uMCP@2		dst:mcp@	src:uMCP@1	
	<====Secondary MPTCP subflow====> <---TCP -->				




Figure 7: Sample Connection Establishment (Implicit Mode)

Subsequent subflows are then sent directly to the MCP (Figure 7). The handling of these subsequent subflows is identical to the one of the explicit mode; only the establishment of the initial subflow differs. Concretely, in reference to Figure 8, once the upstream MCP intercepts an initial subflow, it adds itself to the MPTCP connection by sending ADD_ADDR on the primary subflow. Then, MP_JOIN is sent to the IP address conveyed in ADD_ADDR to create the secondary subflow.

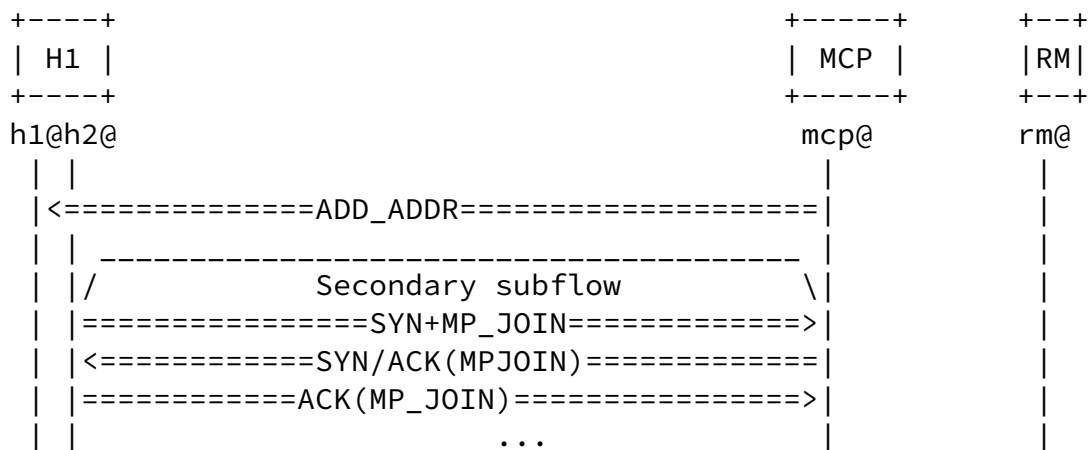


Figure 8: Secondary Subflow Creation (Host-based Implicit Mode)

[5.2.2.1](#). Demux Native MPTCP Flows from Proxied MPTCP Flows

If no explicit signal is included in the initial SYN message, the MCP cannot distinguish "native" MPTCP connections from "proxied" ones. An operator that deploys MCP resources would like to control the MPTCP connections it terminate. For example, an operator may want to enforce a policy that consists in assisting only MPTCP connections that are established by an upstream MCP, not those that are established by an MPTCP host located behind a CPE.

Because MPTCP connections are not destined explicitly to an MCP, an on-path MCP instance will need extra means to distinguish "native" MPTCP connections from "proxied" ones. The subsequent risk is that native MPTCP communications will be reverted to TCP connections as shown in Figure 9 if the MCP is instructed to always relay MPTCP connections to TCP ones. In this example, we suppose that C2 and S2 are MPTCP-capable, but C1 and S1 are not.

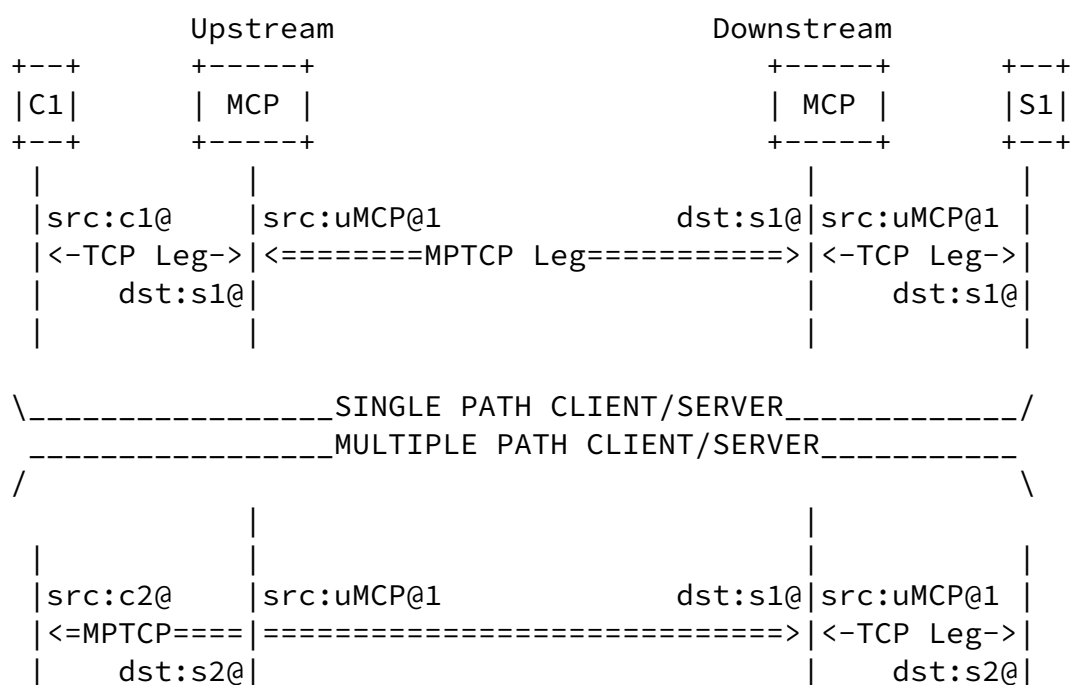


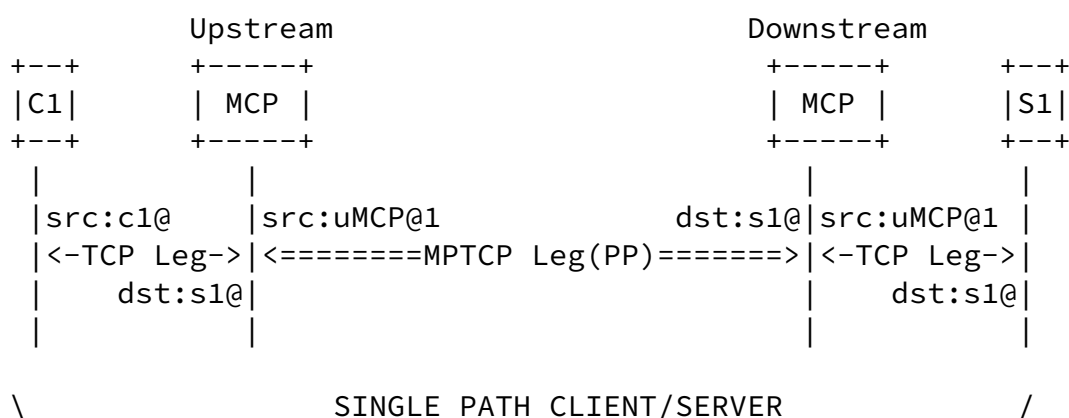


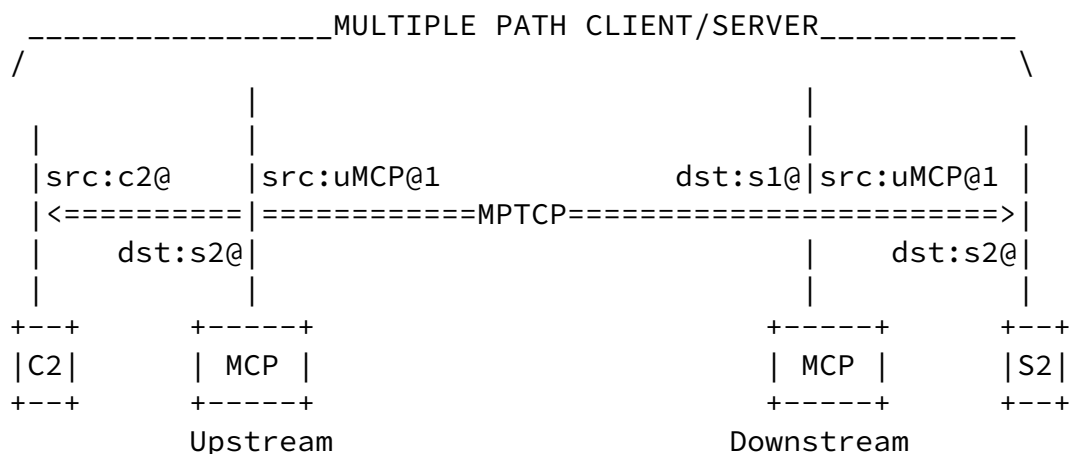
Figure 9: Example of a Broken E2E MPTCP Connection (On-path)

To mitigate this, the upstream MCP may be instructed to insert a `MP_PREFER_PROXY` option only for the MPTCP connections it establishes [[I-D.boucadair-mptcp-plain-mode](#)]. The absence of `MP_PREFER_PROXY` option instances is an explicit indication that this MPTCP connection is a native one. As such, an on-path MCP will not revert this connection into a TCP connection, but will forward packets without any modification to the next hop.

Figure 10 illustrates the results of such procedure: native MPTCP connections are established between MPTCP-capable client and server, while Network-Assisted MPTCP connections are established with the help of MCPs.

Concretely, if the upstream MCP receives a SYN that includes the `MP_PREFER_PROXY` option, it MAY decide to forward it towards its final destination without modifying it. When the downstream MCP receives a SYN that does not include an `MP_PREFER_PROXY`, it forwards it towards its final destination.





Legend:

PP: MP_PREFER_PROXY

Figure 10: Example of a Successful E2E MPTCP Connection (On-path)

5.3. Authorization

The Network Provider that manages the various network attachments (including the MCPs) may enforce authentication and authorization policies using appropriate mechanisms. For example, a non-exhaustive list of methods to achieve authorization is provided hereafter:

- o The network provider may enforce a policy based on the International Mobile Subscriber Identity (IMSI) to verify that a user is allowed to benefit from the aggregation service. If that authorization fails, the Packet Data Protocol (PDP) context /bearer will not be mounted. This method does not require any interaction with the MCP.
- o The network provider may enforce a policy based upon Access Control Lists (ACLs), e.g., at a Broadband Network Gateway (BNG) to control the CPEs that are authorized to communicate with an MCP. These ACLs may be installed as a result of RADIUS exchanges,

for instance ([[I-D.boucadair-mptcp-radius](#)]). This method does not require any interaction with the MCP.

- o The MCP may implement an Ident interface [[RFC1413](#)] to retrieve an

identifier that will be used to assess whether that client is entitled to make use of the aggregation service. Ident exchanges will take place only when receiving the first subflow from a given source IP address.

- o The device that embeds the MCP may also host a RADIUS client that will solicit an AAA server to check whether connections received from a given source IP address are authorized or not ([\[I-D.boucadair-mptcp-radius\]](#)).

A first safeguard against the misuse of MCP resources by illegitimate users (e.g., users with access networks that are not managed by the same service provider that operates the MCP) is to reject MPTCP connections received on the Internet-facing interfaces. Only MPTCP connections received on the customer-facing interfaces of an MCP will be accepted.

Because only the CPE is entitled to establish MPTCP connections with an MCP, ACLs may be installed on the CPE to avoid that internal terminals issue MPTCP connections towards one of the MCPs.

[5.4.](#) MCP Behaviors

The MCP MUST be provided with instructions about the behavior to adopt with regards to the processing of source addresses. The following sub-sections elaborate on various schemes.

[5.4.1.](#) Transparent MCP

Transparent Network-Assisted MPTCP deployment is a deployment where the visible source address of a packet forwarded by an MCP to a remote machine located in the Internet is the IP address of the endhost, not an address that is provisioned to the MCP. In order to intercept incoming traffic, specific IPv4/IPv6 routes are injected so that traffic is redirected towards the MCP.

No dedicated IP address pool is required to the MCP for the Network-Assisted MPTCP service.

[5.4.1.1.](#) IPv4 Address Preservation

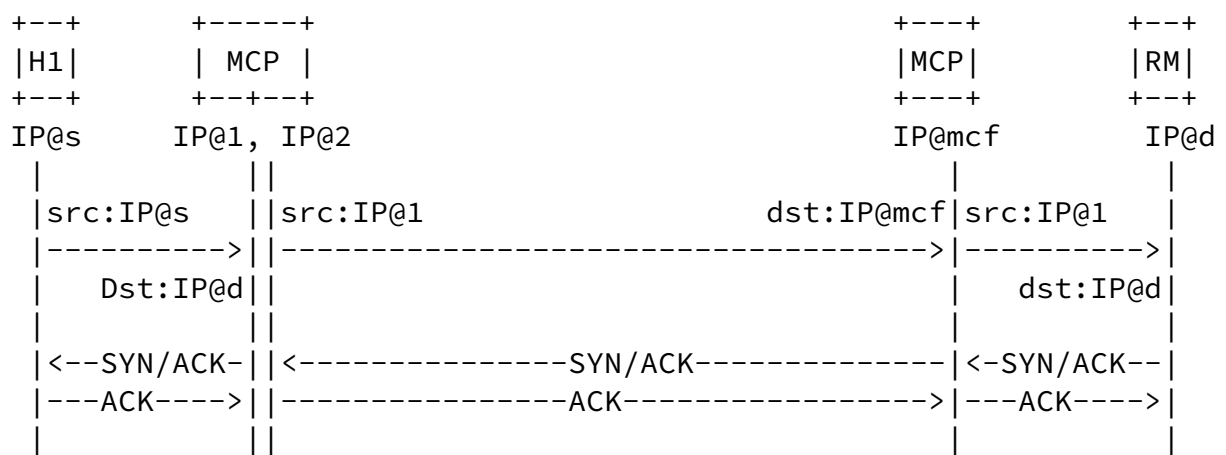
The MCP can be tweaked to behave in the IPv4 address preservation mode. This is the IPv4 address assigned to the endhost (typically,

within a mobile deployment context as discussed in [Section 4.1](#)) or a WAN address of the CPE for the wired case ([Section 4.2](#)).

[5.4.1.2](#). Source IPv6 Prefix Preservation at Network-located MCP

Some IPv6 deployments may require the preservation of the source IPv6 prefix (Figure 11).

This model requires the MCP to support ALGs to accommodate applications with IPv6 address referrals.



Legend:

mcf: MCP Customer-facing Interface

Figure 11: Example of Source IPv6 Prefix Preservation at Network-located MCP (Initial subflow)

[5.4.1.3](#). Source IPv6 Address Preservation at Network-located MCP

Some IPv6 deployments may require the preservation of the source IPv6 address (Figure 12).

This model avoids the need for the MCP to support ALGs to accommodate applications with IPv6 address referrals.

Internet-Draft

Network-Assisted MPTCP

December 2016

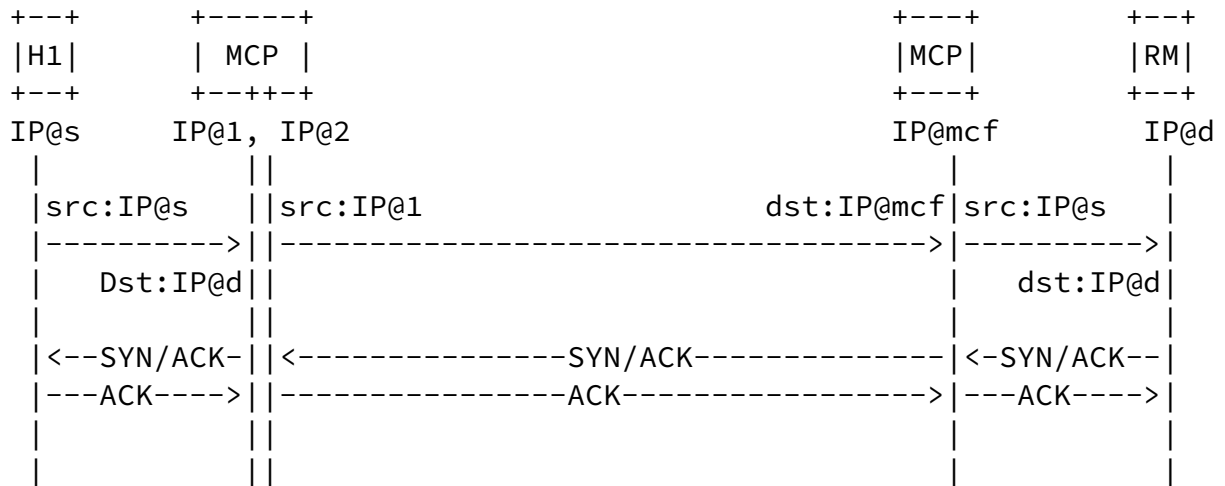


Figure 12: Example of Outgoing SYN with Source Address Preservation

5.4.2. Non-Transparent MCP

Unlike the transport mode, this section focuses on deployments where a dedicated IP address pool is provisioned to the MCP for the Network-Assisted MPTCP service.

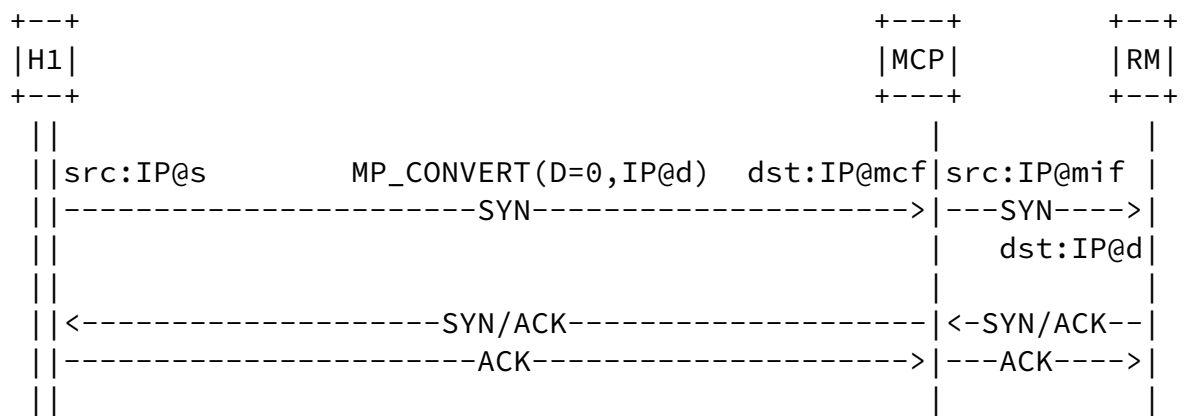
5.4.2.1. IPv4 Address Sharing at the at the Network-located MCP

Because of global IPv4 address depletion, optimization of the IPv4 address usage is mandatory. This obviously includes the IPv4 addresses that are assigned by the MCP at its Internet-facing interfaces (Figure 13 and Figure 14). A pool of global IPv4 addresses is provisioned to the MCP along with possible instructions about the address sharing ratio to apply (see [Appendix B of \[RFC6269\]](#)). Adequate forwarding policies are enforced so that traffic destined to an address of such pool is intercepted by the appropriate MCP.

Internet-Draft

Network-Assisted MPTCP

December 2016



Legend:

mcf: MCP Customer-facing Interface

mif: MCP Internet-facing Interface

Figure 13: Example of Outgoing SYN without Source Address Preservation (Single-ended MCP)

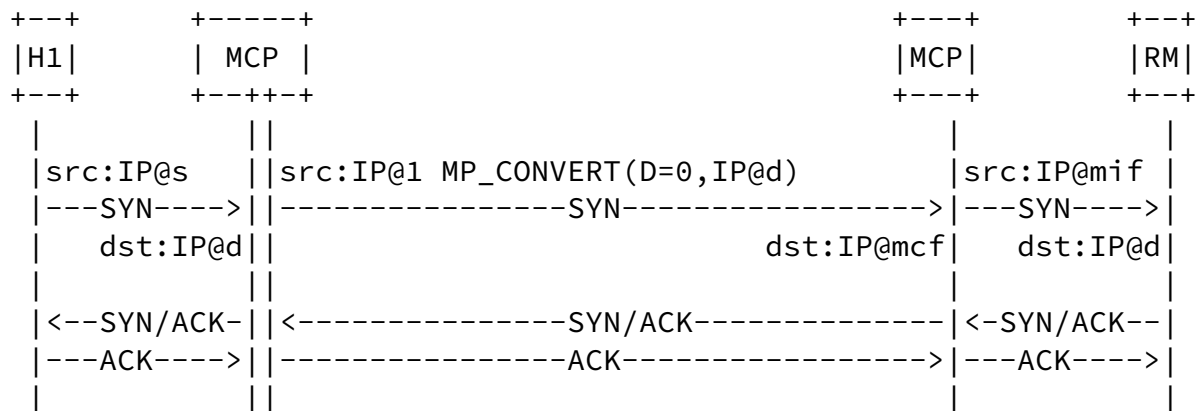


Figure 14: Example of Outgoing SYN without Source Address

[5.4.2.2.](#) IPv4 Address 1:1 Translation at the MCP

For networks that do not face global IPv4 address depletion yet, the MCP can be configured so that source IPv4 addresses of the CPE are replaced with other (public) IPv4 addresses. A pool of global IPv4 addresses is then provisioned to the MCP for this purpose. Rewriting source IPv4 addresses may be used as a means to redirect incoming traffic towards the appropriate MCP.

[5.4.2.3.](#) IPv6 Prefix Sharing (NPTv6) at the Network-located MCP

Rewriting the source IPv6 prefix ([\[RFC6296\]](#)) may be needed to redirect incoming traffic towards the appropriate MCP. A pool of IPv6 prefixes is then provisioned to the MCP for this purpose.

[5.5.](#) Address Family Considerations

Subflows of a given MPTCP connection can be associated to the same address family or may be established with different address families. Also, the Network-Assisted MPTCP using MP_CONVERT IE, regardless of the addressing scheme enforced by each CPE network attachment. In particular, the plain transport mode indifferently accommodates the following combinations.

LAN Leg	MPTCP Legs	TCP Leg towards RM
-----	-----	-----
IPv4	IPv4	IPv4
IPv4	IPv6	IPv4
IPv4	IPv6 & IPv4	IPv4
IPv6	IPv6	IPv6
IPv6	IPv4	IPv6
IPv6	IPv6 & IPv4	IPv6

[5.6.](#) Policies & Configuration Parameters

[5.6.1.](#) Towards End-to-End MPTCP Connections

In order to promote the use of MPTCP end-to-end, the MCP MUST NOT strip MP_CAPABLE from the SYN segments it forwards to their final destination (i.e., server). The communication leg between an MCP and the server may be placed using MPTCP if that server is MPTCP-capable, or falls back to TCP otherwise.

There may be cases where remote servers will not respond to SYN segments with the MP_CAPABLE option, and therefore it is desirable to provide a mechanism to disable relaying MP_CAPABLE to some or all remote hosts.

Also, an MCP SHOULD maintain a cache to record the servers that are known to cause excessive delays. Any subsequent connection to a server that is present in this cache MUST be placed using TCP. Cache entries MUST be flushed out after the expiry of a configurable timer; this timer is set by default to 24h.

This default behavior would lead to the establishment of end-to-end MPTCP connections if the client and servers are both MPTCP-capable with or without the withdrawal of MCPs from the communication.

Whether an MCP must be maintained in the processing of an MPTCP connection that involve MPTCP-capable clients and server is a configurable parameter. The default mode MUST be to maintain the MCP because its presence may solve several complications that may arise in some deployments, such:

1. The use of private IPv4 addresses in some access networks. Typically, additional subflows can not be added to the MPTCP connection without the help of an MCP.
2. The assignment of IPv6 prefix only by some networks. If the server is IPv4-only, subflows that are placed over IPv6 cannot be added to an MPTCP connection towards that server.
3. Some of the access networks are subject to subscription with volume quota.

Figure 15 shows an example of a communication with MCP withdrawal. In this example, the MCP tracks the initial subflow. Upon receipt of

MP_CAPABLE in the SYN/ACK from the server, this connection is not tracked any more by the proxy. Subsequent subflows are handled directly by the endhost and the server.

For cases where some access networks are subject to a volume quota, it is desirable to support a signal to indicate to a remote peer how it must place data into available subflows.

```

-----
/ Host-based Model \
\-----/

```

+-----+	+-----+	+---+
H1	MCP	RM
+-----+	+-----+	+---+
h1@h2@	mcp@	rm@
src:h1	dst:rm@ src:h1	dst:rm@
=====SYN+MP_CAPABLE=====>	=====SYN+MP_CAPABLE=====>	
<=====SYN/ACK+MP_CAPABLE==	<=====SYN/ACK+MP_CAPABLE==	

```

| |=====ACK+MP_CAPABLE=====>|=====ACK+MP_CAPABLE=====>|
| |-----|
| |/      Secondary subflow      \
| |=====SYN+MP_JOIN=====>|
| |<=====SYN/ACK(MPJOIN)=====|
| |=====ACK(MP_JOIN)=====>|
| |
| |      ...
| |

```

Figure 15: Sample Connection Establishment with MCP Withdrawal
(Implicit Mode)

If the MCP is instructed to be involved in the communication even if the server is MPTCP-capable, the flow exchange depicted in Figure 16 will be observed.

```

/-----\
/ Host-based Model \
\-----/

```

+-----+

+-----+

+++

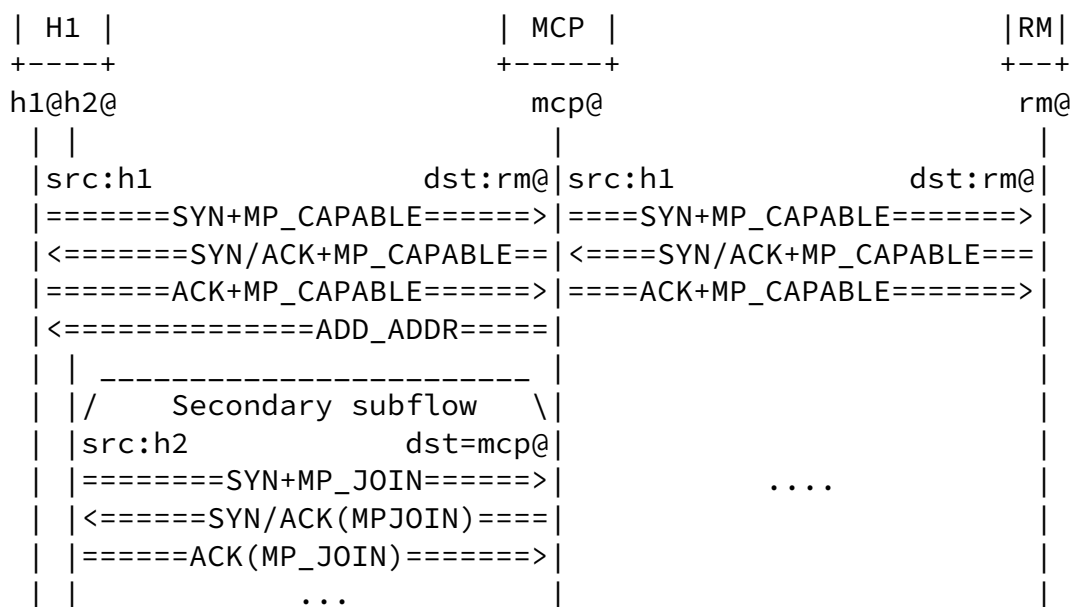


Figure 16: Sample Connection Establishment with MCP Withdrawal (Implicit Mode)

5.6.2. Traffic Distribution Scheme

The logic of traffic distribution over multiple paths is deployment-specific. This document does not require nor preclude any particular traffic distribution scheme. Nevertheless, MCPs MUST be configurable with a parameter to indicate which traffic distribution scheme to enable. Indeed, policies can be enforced by an MCP instance operated by the Network Provider to manage both upstream and downstream traffic. These policies may be subscriber-specific, connection-specific, system-wise, or else.

5.6.3. Flows Eligible to Multipath Service

The Multipath Client and MCPs may be provided with a set of classification policies to help electing flows for the MPTCP service. These policies may be provisioned either statically and dynamically (or a combination thereof).

Also, multiple MCPs may serve a given end-user, as a function of the nature of the service or the traffic to be forwarded over MPTCP connections. For example, an MCP may be used by a service provider to proceed with CPE-targeted maintenance operations, whereas another

MCP may be configured to service multi-path communications initiated by a set of end-users.

[5.6.4.](#) TCP Fragmentation

Methods to avoid TCP fragmentation, such as rewriting the TCP Maximum Segment Size (MSS) option, must be supported by MCPs.

[5.6.5.](#) DSCP Preservation

The MCP MAY be configured to preserve the same DSCP marking or enforce DSCP re-marking policies. DSCP preservation MUST be enabled by default.

[5.6.6.](#) Supported Transport Protocols

The MCP supports TCP by design. Additional transport protocols SHOULD be supported. A configuration parameter MUST be supported by the MCP to indicate which transport protocols can be relayed into an MPTCP connection.

[5.6.7.](#) Logging

If the MCP is used in global IPv4 address sharing environments, the logging recommendations discussed in [Section 4 of \[RFC6888\]](#) need to be considered. Security-related issues encountered in address sharing environments are documented in [Section 13 of \[RFC6269\]](#). A configuration parameter should be supported to enable/disable the logging function.

[6.](#) IANA Considerations

This document does not request any action from IANA.

[7.](#) Security Considerations

MPTCP-related security threats are discussed in [\[RFC6181\]](#) and [\[RFC6824\]](#). Additional considerations are discussed in the following sub-sections.

[7.1.](#) Privacy

The MCP may have access to privacy-related information (e.g., IMSI, link identifier, subscriber credentials, etc.). The MCP MUST NOT leak such sensitive information outside a local domain.

[7.2.](#) Denial-of-Service (DoS)

Means to protect the MCP against Denial-of-Service (DoS) attacks MUST be enabled. Such means include the enforcement of ingress filtering policies at the network boundaries [[RFC2827](#)].

In order to prevent the exhaustion of MCP's resources by establishing a large number of simultaneous subflows for each MPTCP connection, the MCP administrator SHOULD limit the number of allowed subflows per CPE for a given connection. Means to protect against SYN flooding attacks MUST also be enabled ([[RFC4987](#)]).

Attacks that originate outside of the domain can be prevented if ingress filtering policies are enforced. Nevertheless, attacks from within the network between a host and an MCP instance are yet another actual threat. Means to ensure that illegitimate nodes cannot connect to a network should be implemented.

[7.3.](#) Illegitimate MCP

Traffic theft is a risk if an illegitimate MCP is inserted in the path. Indeed, inserting an illegitimate MCP in the forwarding path allows traffic intercept and can therefore provide access to sensitive data issued by or destined to a host. To mitigate this threat, secure means to discover an MCP should be enabled.

[7.4.](#) High Rate Reassembly

The MCP may perform packet reassembly. Some security-related issues are discussed in [[RFC4963](#)][[RFC1858](#)][[RFC3128](#)].

[8.](#) Contributors

The following individuals contributed to this document:

Denis Behaghel
OneAccess

Email: Denis.Behaghel@oneaccess-net.com

Stefano Secci

Universite Pierre et Marie Curie
Paris
France

Email: stefano.secci@lip6.fr

Boucadair, et al.

Expires June 10, 2017

[Page 24]

Internet-Draft

Network-Assisted MPTCP

December 2016

Suresh Vinapamula
Juniper
1137 Innovation Way
Sunnyvale, CA 94089
USA

Email: Sureshk@juniper.net

SungHoon Seo
Korea Telecom
Seoul
Korea

Email: sh.seo@kt.com

Wouter Cloetens
SoftAtHome
Vaartdijk 3 701
3018 Wijkmaal
Belgium

Email: wouter.cloetens@softathome.com

Ullrich Meyer
Vodafone
Germany

Email: ullrich.meyer@vodafone.com

Luis M. Contreras
Telefonica

Spain

Email: luismiguel.contrerasmurillo@telefonica.com

Bart Peirens
Proximus

Email: bart.peirens@proximus.com

Boucadair, et al.

Expires June 10, 2017

[Page 25]

Internet-Draft

Network-Assisted MPTCP

December 2016

[9.](#) Acknowledgements

Many thanks to Chi Dung Phung, Mingui Zhang, Rao Shoaib, Yoshifumi Nishida, and Christoph Paasch for the comments.

Thanks to Ian Farrer, Mikael Abrahamsson, Alan Ford, Dan Wing, and Sri Gundavelli for the fruitful discussions held during the IETF#95 meeting.

Special thanks to Pierrick Seite, Yannick Le Goff, Fred Klamm, and Xavier Grall for their valuable comments.

Thanks also to Olaf Schleusing, Martin Gysi, Thomas Zasowski, Andreas Burkhard, Silka Simmen, Sandro Berger, Michael Melloul, Jean-Yves Flahaut, Adrien Desportes, Gregory Detal, Benjamin David, Arun Srinivasan, and Raghavendra Mallya for the discussion.

[10.](#) References

[10.1.](#) Normative References

[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, "An MPTCP Option for Network-Assisted MPTCP", [draft-boucadair-mptcp-plain-mode-09](#) (work in progress), October 2016.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", [RFC 6824](#), DOI 10.17487/RFC6824, January 2013, <<http://www.rfc-editor.org/info/rfc6824>>.

10.2. Informative References

- [I-D.boucadair-mptcp-dhc]
Boucadair, M., Jacquenet, C., and T. Reddy, "DHCP Options for Network-Assisted Multipath TCP (MPTCP)", [draft-boucadair-mptcp-dhc-06](#) (work in progress), October 2016.

Boucadair, et al. Expires June 10, 2017 [Page 26]

Internet-Draft Network-Assisted MPTCP December 2016

- [I-D.boucadair-mptcp-radius]
Boucadair, M. and C. Jacquenet, "RADIUS Extensions for Network-Assisted Multipath TCP (MPTCP)", [draft-boucadair-mptcp-radius-03](#) (work in progress), October 2016.
- [I-D.ietf-mptcp-experience]
Bonaventure, O., Paasch, C., and G. Detal, "Use Cases and Operational Experience with Multipath TCP", [draft-ietf-mptcp-experience-07](#) (work in progress), October 2016.
- [I-D.ietf-mptcp-rfc6824bis]
Ford, A., Raiciu, C., Handley, M., Bonaventure, O., and C. Paasch, "TCP Extensions for Multipath Operation with Multiple Addresses", [draft-ietf-mptcp-rfc6824bis-07](#) (work in progress), October 2016.
- [RFC1413] St. Johns, M., "Identification Protocol", [RFC 1413](#), DOI 10.17487/RFC1413, February 1993, <<http://www.rfc-editor.org/info/rfc1413>>.
- [RFC1858] Ziemba, G., Reed, D., and P. Traina, "Security

Considerations for IP Fragment Filtering", [RFC 1858](#), DOI 10.17487/RFC1858, October 1995, <<http://www.rfc-editor.org/info/rfc1858>>.

- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", [BCP 38](#), [RFC 2827](#), DOI 10.17487/RFC2827, May 2000, <<http://www.rfc-editor.org/info/rfc2827>>.
- [RFC3128] Miller, I., "Protection Against a Variant of the Tiny Fragment Attack ([RFC 1858](#))", [RFC 3128](#), DOI 10.17487/RFC3128, June 2001, <<http://www.rfc-editor.org/info/rfc3128>>.
- [RFC4963] Heffner, J., Mathis, M., and B. Chandler, "IPv4 Reassembly Errors at High Data Rates", [RFC 4963](#), DOI 10.17487/RFC4963, July 2007, <<http://www.rfc-editor.org/info/rfc4963>>.
- [RFC4987] Eddy, W., "TCP SYN Flooding Attacks and Common Mitigations", [RFC 4987](#), DOI 10.17487/RFC4987, August 2007, <<http://www.rfc-editor.org/info/rfc4987>>.
- [RFC6181] Bagnulo, M., "Threat Analysis for TCP Extensions for Multipath Operation with Multiple Addresses", [RFC 6181](#), DOI 10.17487/RFC6181, March 2011, <<http://www.rfc-editor.org/info/rfc6181>>.

- [RFC6269] Ford, M., Ed., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", [RFC 6269](#), DOI 10.17487/RFC6269, June 2011, <<http://www.rfc-editor.org/info/rfc6269>>.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", [RFC 6296](#), DOI 10.17487/RFC6296, June 2011, <<http://www.rfc-editor.org/info/rfc6296>>.
- [RFC6888] Perreault, S., Ed., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", [BCP 127](#), [RFC 6888](#), DOI 10.17487/RFC6888, April 2013, <<http://www.rfc-editor.org/info/rfc6888>>.

[RFC7974] Williams, B., Boucadair, M., and D. Wing, "An Experimental TCP Option for Host Identification", [RFC 7974](#), DOI 10.17487/RFC7974, October 2016, <<http://www.rfc-editor.org/info/rfc7974>>.

Authors' Addresses

Mohamed Boucadair (editor)
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Christian Jacquenet (editor)
Orange
Rennes
France

Email: christian.jacquenet@orange.com

Olivier Bonaventure (editor)
Tessares
Belgium

Email: olivier.bonaventure@tessares.net

Wim Henderickx (editor)
Nokia/Alcatel-Lucent
Belgium

Email: wim.henderickx@alcatel-lucent.com

Robert Skog (editor)
Ericsson

Email: robert.skog@ericsson.com