

Network Working Group  
Internet-Draft  
Intended status: Experimental  
Expires: September 5, 2015

M. Boucadair  
C. Jacquenet  
France Telecom  
March 4, 2015

MPTCP Connectivity Checks  
draft-boucadair-mptcp-connectivity-checks-00

## Abstract

This document specifies an extension to minimize the delay induced by the presence of MPTCP-unfriendly nodes in some of the paths selected by a MPTCP endpoint, and which may support the establishment of MPTCP subflows. Concretely, this procedure allows a MPTCP endpoint to assess whether the networks the endpoint connects to are compliant with MPTCP signals or not. The procedure is not enabled for every new MPTCP connection; it is activated upon bootstrap or when a network attachment change occurs (e.g., attach to a new network, discover a new external IP address, etc.).

## 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-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 September 5, 2015.

Internet-Draft

MPTCP Compatibility Assessement

March 2015

## Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

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

## Table of Contents

<a href="#">1.</a>	Introduction . . . . .	<a href="#">2</a>
<a href="#">2.</a>	Problem Space . . . . .	<a href="#">2</a>
<a href="#">3.</a>	Proposed Solution . . . . .	<a href="#">4</a>
<a href="#">4.</a>	Security Considerations . . . . .	<a href="#">7</a>
<a href="#">5.</a>	IANA Considerations . . . . .	<a href="#">7</a>
<a href="#">6.</a>	Acknowledgements . . . . .	<a href="#">7</a>
<a href="#">7.</a>	References . . . . .	<a href="#">7</a>
<a href="#">7.1.</a>	Normative References . . . . .	<a href="#">7</a>
<a href="#">7.2.</a>	Informative References . . . . .	<a href="#">7</a>
	Authors' Addresses . . . . .	<a href="#">8</a>

## [1.](#) Introduction

This document specifies an extension that minimizes the delay induced by the presence of MPTCP-unfriendly nodes in the some of the paths selected by a MPTCP endpoint, and which may support the establishment of MPTCP subflows.

The problem space is further described in [Section 2](#), while a proposed solution is discussed in [Section 3](#).

## [2.](#) Problem Space


Advanced flow-aware service functions (e.g., Performance Enhancement Proxies ([[RFC3135](#)]), NATs [[RFC3022](#)], CGNs [[RFC6888](#)], DS-Lite AFTR [[RFC6333](#)], NAT64 [[RFC6146](#)], NPTv6 [[RFC6296](#)], firewalls, etc.) are

required to achieve various objectives such as IP address sharing, firewalling, avoid covert channels, detect and protect against ever increasing DDoS attacks, etc.

Removing those functions is not an option because they are used to address constraints that are often typical of the current yet protean Internet situation (global IPv4 address depletion comes to mind, but also the plethora of services with different QoS/security/robustness requirements, etc.), and this is even exacerbated by environment-specific designs (e.g., the nature and the number of service functions that need to be invoked at the Gi interface of a mobile infrastructure). Moreover, these sophisticated service functions are located in the network but also in service platforms, or intermediate entities (e.g., CDNs). A flow-aware device can be embedded in a CPE (Customer Premises Equipment) and/or hosted in the network provider's side.

In the meantime, the presence of these flow-aware functions complicates the introduction of new protocols or the introduction of additional features for existing ones. Also, because some of these flow-aware functions do not expose a deterministic behavior, additional complications are encountered even if the protocol design includes built-in features to detect (and possibly accommodate) the presence of such functions. This document focuses on MPTCP [[RFC6824](#)].

MPTCP supports a mechanism to fallback to TCP when a flow-aware function interferes with MPTCP signals. Figure 1 and Figure 2 show two typical examples of the fallback behavior. It is out of scope of this document to list all the possible fallback scenarios. Refer to [[RFC6824](#)] for more details about the exact fallback behavior.






Figure 1: Fallback Case 1

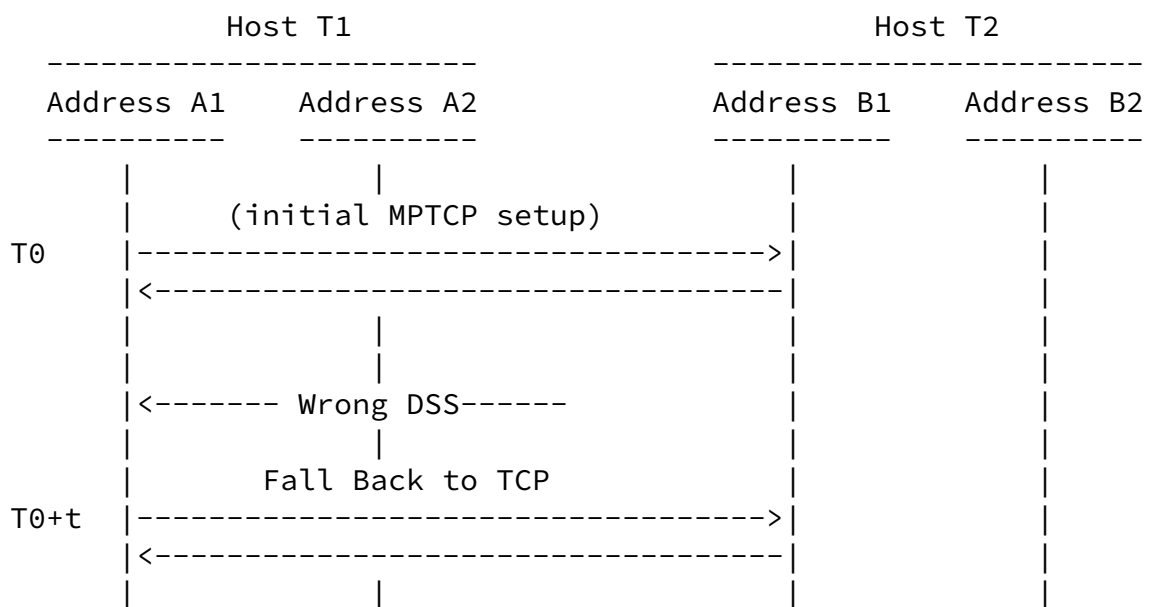


Figure 2: Fallback Case 2

This behavior, if adopted for every new connection, will have negative impacts on the quality of experience as perceived by a user. This is not desirable.

### 3. Proposed Solution

The problem in [Section 2](#) can be originated by MPTCP-unfriendly service function(s) located at the initiator side, the receiver side, or the network in between. In order to avoid increasing the delay to establish a TCP-based connection involving an MPTCP-enabled endpoint,

this document suggests the use of connectivity checks to assess whether available paths as perceived by an MPTCP-enabled endpoint can safely convey MPTCP signals. Doing so, the results of such connectivity checks allow an MPTCP-enabled endpoint to decide whether MPTCP needs to be disabled for all or part of its available network attachments. This also allows the endpoint to identify which paths cannot be used to establish the first subflow. Network attachments that aren't MPTCP-friendly are tagged as such.

A dedicated functional element (referred to as MPTCP Connectivity Check Server (MCCS)) is proposed to assess the MPTCP friendliness of its network attachments. MCCS is attached to a network that does not involve any MPTCP-unfriendly service function. If any MPTCP problem is encountered when establishing a connection with an MCCS, it is an indication that the issue is located at the remote MPTCP-enabled endpoint.

This procedure is activated upon bootstrap or when a network attachment change occurs (e.g., attach to a new network, discover a

new external IP address, etc.); it is not executed for every new MPTCP connection:

- o An MPTCP-enabled endpoint is configured with one or a list of MCCS servers.

A well-known name can be used for this purpose. The name is then passed to the name resolution library (e.g., [Section 6.1.1 of \[RFC1123\]](#)) to retrieve the corresponding IP address(es) (IPv4, IPv6 or both).

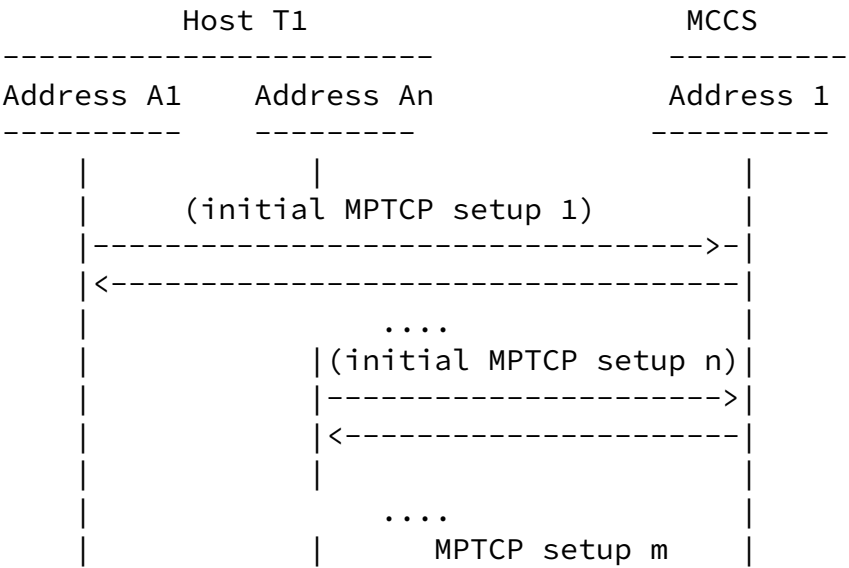
- o The MPTCP-enabled terminal then establishes MPTCP connections based upon all the IP addresses that have been retrieved (or a subset thereof). Executing the procedure with more than one MCCS (if available) is RECOMMENDED.

These MPTCP connections are meant to assess for each available network attachment, interface, discovered external IP address, etc., that the path that will be solicited when issuing packets does not break MPTCP connections.

The tests are executed both from the MPTCP-enabled endpoint and

also the MCCS. The extension defined [\[I-D.boucadair-mptcp-symmetric\]](#) is RECOMMENDED so that the MCCS can initiate MPTCP connections, add new subflows to an active MPTCP connection, etc.

The tests are performed to cover all failure cases (e.g., strip MPTCP signals, alter some options, corrupted DSS, etc.). An example is depicted in Figure 3 for illustration purposes.






Figure 3: An example of connectivity checks

- o As an outcome of the previous step, the MPTCP-enabled endpoint can conclude whether all or some of its available network attachments can "safely" be used when establishing MPTCP connections.

Interfaces/address/networks that are MPTCP-unfriendly are tagged accordingly; those are not used when establishing a new MPTCP connection with a remote peer or to add a new subflow to an existing MPTCP connection.

In particular, an MPTCP-enabled endpoint will not use MPTCP when all its network attachments are MPTCP-unfriendly. This covers in particular the situation where the endpoint initialed the connection or when the MPTCP device receives an MPTCP connection (e.g., user-generated content context).

- o This procedure is executed whenever network conditions change, as perceived by an MPTCP-enabled endpoint.

In the context of [[I-D.deng-mptcp-proxy](#)], this procedure can be implemented at the CPE side on behalf of the MPTCP-enabled terminals the CPE is connected to in the user premises. An MPTCP-enabled terminal located behind a CPE assumes by default that its default gateway is its MCCS. Doing so, this design allows to avoid re-using the same connectivity checks for all the terminals located behind a CPE.

A deployment option would consist in integrating connectivity check capabilities in STUN servers [[RFC5389](#)]; an extension would be needed for that purpose, though.

#### [4.](#) Security Considerations

MPTCP-related security considerations are documented in [[RFC6824](#)] and [[I-D.ietf-mptcp-attacks](#)].

Security-consideration specific to MCCA will be discussed.

## [5.](#) IANA Considerations

This document does not require any action from IANA.

## [6.](#) Acknowledgements

TBC

## [7.](#) References

### [7.1.](#) Normative References

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

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

### [7.2.](#) Informative References

[I-D.boucadair-mptcp-symmetric]  
Boucadair, M. and C. Jacquenet, "An Extension to MPTCP for Symmetrical Sub-Flow Management", March 2015,  
<<https://tools.ietf.org/html/draft-boucadair-mptcp-symmetric-01>>.

[I-D.deng-mptcp-proxy]  
Lingli, D., Liu, D., Sun, T., Boucadair, M., and G. Cauchie, "Use-cases and Requirements for MPTCP Proxy in ISP Networks", [draft-deng-mptcp-proxy-01](#) (work in progress), October 2014.



- Bagnulo, M., Paasch, C., Gont, F., Bonaventure, O., and C. Raiciu, "Analysis of MPTCP residual threats and possible fixes", [draft-ietf-mptcp-attacks-03](#) (work in progress), February 2015.
- [RFC1123] Braden, R., "Requirements for Internet Hosts - Application and Support", STD 3, [RFC 1123](#), October 1989.
- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", [RFC 3022](#), January 2001.
- [RFC3135] Border, J., Kojo, M., Griner, J., Montenegro, G., and Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations", [RFC 3135](#), June 2001.
- [RFC5389] Rosenberg, J., Mahy, R., Matthews, P., and D. Wing, "Session Traversal Utilities for NAT (STUN)", [RFC 5389](#), October 2008.
- [RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", [RFC 6146](#), April 2011.
- [RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", [RFC 6296](#), June 2011.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", [RFC 6333](#), August 2011.
- [RFC6888] Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", [BCP 127](#), [RFC 6888](#), April 2013.

#### Authors' Addresses

Mohamed Boucadair  
France Telecom  
Rennes 35000  
France

Email: [mohamed.boucadair@orange.com](mailto:mohamed.boucadair@orange.com)

Christian Jacquenet  
France Telecom  
Rennes 35000  
France

Email: [christian.jacquenet@orange.com](mailto:christian.jacquenet@orange.com)

