

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: May 26, 2011

M. Scharf
Alcatel-Lucent Bell Labs
A. Ford
Roke Manor Research
November 22, 2010

MPTCP Application Interface Considerations
draft-scharf-mptcp-api-04

Abstract

Multipath TCP (MPTCP) adds the capability of using multiple paths to a regular TCP session. Even though it is designed to be totally backward compatible to applications, the data transport differs compared to regular TCP, and there are several additional degrees of freedom that applications may wish to exploit. This document summarizes the impact that MPTCP may have on applications, such as changes in performance. Furthermore, it discusses compatibility issues of MPTCP in combination with non-MPTCP-aware applications. Finally, the document describes a basic application interface for MPTCP-aware applications that provides access to multipath address information and a level of control equivalent to regular TCP.

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 May 26, 2011.

Copyright Notice

Copyright (c) 2010 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

- [1. Introduction](#) [4](#)
- [2. Terminology](#) [5](#)
- [3. Comparison of MPTCP and Regular TCP](#) [5](#)
 - [3.1. Performance Impact](#) [6](#)
 - [3.1.1. Throughput](#) [6](#)
 - [3.1.2. Delay](#) [6](#)
 - [3.1.3. Resilience](#) [7](#)
 - [3.2. Potential Problems](#) [7](#)
 - [3.2.1. Impact of Middleboxes](#) [7](#)
 - [3.2.2. Outdated Implicit Assumptions](#) [8](#)
 - [3.2.3. Security Implications](#) [8](#)
- [4. Operation of MPTCP with Legacy Applications](#) [8](#)
 - [4.1. Overview of the MPTCP Network Stack](#) [8](#)
 - [4.2. Address Issues](#) [9](#)
 - [4.2.1. Specification of Addresses by Applications](#) [9](#)
 - [4.2.2. Querying of Addresses by Applications](#) [10](#)
 - [4.3. Socket Option Issues](#) [11](#)
 - [4.3.1. General Guideline](#) [11](#)
 - [4.3.2. Disabling of the Nagle Algorithm](#) [11](#)
 - [4.3.3. Buffer Sizing](#) [11](#)
 - [4.3.4. Other Socket Options](#) [12](#)
 - [4.4. Default Enabling of MPTCP](#) [12](#)
 - [4.5. Summary of Advices to Application Developers](#) [12](#)
- [5. Basic API for MPTCP-aware Applications](#) [13](#)
 - [5.1. Design Considerations](#) [13](#)
 - [5.2. Requirements on the Basic MPTCP API](#) [13](#)
 - [5.3. Sockets Interface Extensions by the Basic MPTCP API](#) [15](#)
 - [5.3.1. Overview](#) [15](#)
 - [5.3.2. Enabling and Disabling of MPTCP](#) [16](#)
 - [5.3.3. Binding MPTCP to Specified Addresses](#) [17](#)
 - [5.3.4. Querying the MPTCP Subflow Addresses](#) [17](#)
 - [5.3.5. Getting a Unique Connection Identifier](#) [18](#)
- [6. Other Compatibility Issues](#) [18](#)
 - [6.1. Usage of the SCTP Socket API](#) [18](#)
 - [6.2. Incompatibilities with other Multihoming Solutions](#) [18](#)
 - [6.3. Interactions with DNS](#) [19](#)
- [7. Security Considerations](#) [19](#)
- [8. IANA Considerations](#) [19](#)

- [9. Conclusion](#) [19](#)
- [10. Acknowledgments](#) [20](#)
- [11. References](#) [20](#)
 - [11.1. Normative References](#) [20](#)
 - [11.2. Informative References](#) [21](#)
- [Appendix A. Requirements on a Future Advanced MPTCP API](#) [21](#)
 - [A.1. Design Considerations](#) [21](#)
 - [A.2. MPTCP Usage Scenarios and Application Requirements](#) [22](#)
 - [A.3. Potential Requirements on an Advanced MPTCP API](#) [24](#)
- [Appendix B. Change History of the Document](#) [25](#)

1. Introduction

Multipath TCP adds the capability of using multiple paths to a regular TCP session [1]. The motivations for this extension include increasing throughput, overall resource utilisation, and resilience to network failure, and these motivations are discussed, along with high-level design decisions, as part of the Multipath TCP architecture [4]. The MPTCP protocol [5] offers the same reliable, in-order, byte-stream transport as TCP, and is designed to be backward compatible with both applications and the network layer. It requires support inside the network stack of both endpoints.

This document first presents the impacts that MPTCP may have on applications, such as performance changes compared to regular TCP. Second, it defines the interoperation of MPTCP and applications that are unaware of the multipath transport. MPTCP is designed to be usable without any application changes, but some compatibility issues have to be taken into account. Third, this memo specifies a basic Application Programming Interface (API) for MPTCP-aware applications. The API presented here is an extension to the regular TCP API to allow an MPTCP-aware application the equivalent level of control and access to information of an MPTCP connection that would be possible with the standard TCP API on a regular TCP connection.

An advanced API for MPTCP is outside the scope of this document. Such an advanced API could offer a more fine-grained control over multipath transport functions and policies. The appendix includes a brief, non-compulsory list of potential features of such an advanced API.

The de facto standard API for TCP/IP applications is the "sockets" interface. This document provides an abstract definition of MPTCP-specific extensions to this interface. These are operations that can be used by an application to get or set additional MPTCP-specific information on a socket, in order to provide an equivalent level of information and control over MPTCP as exists for an application using regular TCP. It is up to the applications, high-level programming languages, or libraries to decide whether to use these optional extensions. For instance, an application may want to turn on or off the MPTCP mechanism for certain data transfers, or limit its use to certain interfaces. The abstract specification is in line with the Posix standard [8] as much as possible.

There are also various related extensions of the sockets interface: [12] specifies sockets API extensions for a multihoming shim layer. The API enables interactions between applications and the multihoming shim layer for advanced locator management and for access to information about failure detection and path exploration.

Experimental extensions to the sockets API are also defined for the Host Identity Protocol (HIP) [13] in order to manage the bindings of identifiers and locator. Further related API extensions exist for IPv6 [10], Mobile IP [11], and SCTP [14]. There can be interactions or incompatibilities of these APIs with MPTCP, which are discussed later in this document.

Some network stack implementations, specially on mobile devices, have centralized connection managers or other higher-level APIs to solve multi-interface issues, as surveyed in [16]. Their interaction with MPTCP is outside the scope of this note.

The target readers of this document are application developers whose software may benefit significantly from MPTCP. This document also provides the necessary information for developers of MPTCP to implement the API in a TCP/IP network stack.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [3].

This document uses the MPTCP terminology introduced in [5].

Concerning the API towards applications, the following terms are distinguished:

- o Legacy API: The interface towards TCP that is currently used by applications. This document explains the impact of MPTCP for such applications, as well as resulting issues.
- o Basic API: A simple extension of TCP's interface for applications that are aware of MPTCP. This document abstractly describes this interface, which provides access to multipath address information and a level of control equivalent to regular TCP.
- o Advanced API: An API that offers more fine-grained control over the MPTCP behaviour. Its detailed specification is outside scope of this document.

3. Comparison of MPTCP and Regular TCP

This section discusses the impact that the use of MPTCP will have on applications, in comparison to what may be expected from the use of regular TCP.

3.1. Performance Impact

One of the key goals of adding multipath capability to TCP is to improve the performance of a transport connection by load distribution over separate subflows across potentially disjoint paths. Furthermore, it is an explicit goal of MPTCP that it should not provide a worse performing connection that would have existed through the use of single-path TCP. A corresponding congestion control algorithm is described in [7]. The following sections summarize the performance impact of MPTCP as seen by an application.

3.1.1. Throughput

The most obvious performance improvement that will be gained with the use of MPTCP is an increase in throughput, since MPTCP will pool more than one path (where available) between two endpoints. This will provide greater bandwidth for an application. If there are shared bottlenecks between the flows, then the congestion control algorithms will ensure that load is evenly spread amongst regular and multipath TCP sessions, so that no end user receives worse performance than single-path TCP.

This performance increase additionally means that an MPTCP session could achieve throughput that is greater than the capacity of a single interface on the device. If any applications make assumptions about interfaces due to throughput (or vice versa), they must take this into account (although an MPTCP implementation must always respect an application's request for a particular interface).

Furthermore, the flexibility of MPTCP to add and remove subflows as paths change availability could lead to a greater variation, and more frequent change, in connection bandwidth. Applications that adapt to available bandwidth (such as video and audio streaming) may need to adjust some of their assumptions to most effectively take this into account.

The transport of MPTCP signaling information results in a small overhead. If multiple subflows share a same bottleneck, this overhead slightly reduces the capacity that is available for data transport. Yet, this potential reduction of throughput will be neglectible in many usage scenarios, and the protocol contains optimisations in its design so that this overhead is minimal.

3.1.2. Delay

If the delays on the constituent subflows of an MPTCP connection differ, the jitter perceivable to an application may appear higher as the data is spread across the subflows. Although MPTCP will ensure

in-order delivery to the application, the application must be able to cope with the data delivery being burstier than may be usual with single-path TCP. Since burstiness is commonplace on the Internet today, it is unlikely that applications will suffer from such an impact on the traffic profile, but application authors may wish to consider this in future development.

In addition, applications that make round trip time (RTT) estimates at the application level may have some issues. Whilst the average delay calculated will be accurate, whether this is useful for an application will depend on what it requires this information for. If a new application wishes to derive such information, it should consider how multiple subflows may affect its measurements, and thus how it may wish to respond. In such a case, an application may wish to express its scheduling preferences, as described later in this document.

[3.1.3.](#) Resilience

The use of multiple subflows simultaneously means that, if one should fail, all traffic will move to the remaining subflow(s), and additionally any lost packets can be retransmitted on these subflows.

Subflow failure may be caused by issues within the network, which an application would be unaware of, or interface failure on the node. An application may, under certain circumstances, be in a position to be aware of such failure (e.g. by radio signal strength, or simply an interface enabled flag), and so must not make assumptions of an MPTCP flow's stability based on this. An MPTCP implementation must never override an application's request for a given interface, however, so the cases where this issue may be applicable are limited.

[3.2.](#) Potential Problems

[3.2.1.](#) Impact of Middleboxes

MPTCP has been designed in order to pass through the majority of middleboxes. Empirical evidence suggests that new TCP options can successfully be used on most paths in the Internet. Nevertheless some middleboxes may still refuse to pass MPTCP messages due to the presence of TCP options, or they may strip TCP options. If this is the case, MPTCP should fall back to regular TCP. Although this will not create a problem for the application (its communication will be set up either way), there may be additional (and indeed, user-perceivable) delay while the first handshake fails. A detailed discussion of the various fallback mechanisms, for failures occurring at different points in the connection, is presented in [\[5\]](#).

There may also be middleboxes that transparently change the length of content. If such middleboxes are present, MPTCP's reassembly of the byte stream in the receiver is difficult. Still, MPTCP can detect such middleboxes and then fall back to regular TCP. An overview of the impact of middleboxes is presented in [4] and MPTCP's mechanisms to work around these are presented and discussed in [5].

MPTCP can also have other unexpected implications. For instance, intrusion detection systems could be triggered. A full analysis of MPTCP's impact on such middleboxes is for further study after deployment experiments.

3.2.2. Outdated Implicit Assumptions

In regular TCP, there is a one-to-one mapping of the socket interface to a flow through a network. Since MPTCP can make use of multiple subflows, applications cannot implicitly rely on this one-to-one mapping any more. Applications that require the transport along a single path can disable the use of MPTCP as described later in this document. Examples include monitoring tools that want to measure the available bandwidth on a path, or routing protocols such as BGP that require the use of a specific link.

3.2.3. Security Implications

The support for multiple IP addresses within one MPTCP connection can result in additional security vulnerabilities, such as possibilities for attackers to hijack connections. The protocol design of MPTCP minimizes this risk. An attacker on one of the paths can cause harm, but this is hardly an additional security risk compared to single-path TCP, which is vulnerable to man-in-the-middle attacks, too. A detailed thread analysis of MPTCP is published in [6].

4. Operation of MPTCP with Legacy Applications

4.1. Overview of the MPTCP Network Stack

MPTCP is an extension of TCP, but it is designed to be backward compatible for legacy applications. TCP interacts with other parts of the network stack by different interfaces. The de facto standard API between TCP and applications is the sockets interface. The position of MPTCP in the protocol stack can be illustrated in Figure 1.

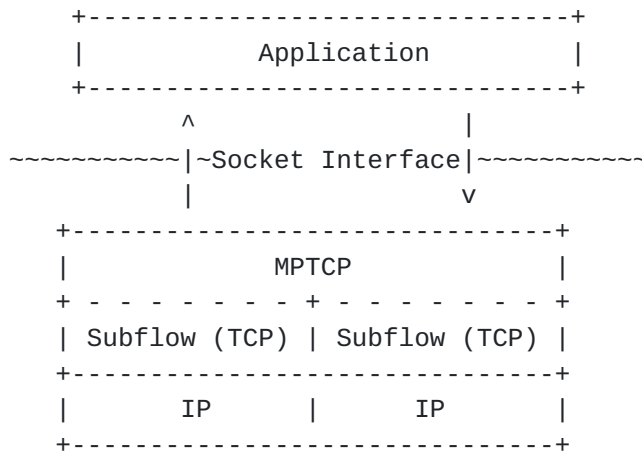


Figure 1: MPTCP protocol stack

In general, MPTCP can affect all interfaces that make assumptions about the coupling of a TCP connection to a single IP address and TCP port pair, to one sockets endpoint, to one network interface, or to a given path through the network.

This means that there are two classes of applications:

- o Legacy applications: These applications are unaware of MPTCP and use the existing API towards TCP without any changes. This is the default case.
- o MPTCP-aware applications: These applications indicate support for an enhance MPTCP interface. This document specified a minimum set of API extensions for such applications.

In the following, it is discussed to which extent MPTCP affects legacy applications using the existing sockets API. The existing sockets API implies that applications deal with data structures that store, amongst others, the IP addresses and TCP port numbers of a TCP connection. A design objective of MPTCP is that legacy applications can continue to use the established sockets API without any changes. However, in MPTCP there is a one-to-many mapping between the socket endpoint and the subflows. This has several subtle implications for legacy applications using sockets API functions.

4.2. Address Issues

4.2.1. Specification of Addresses by Applications

During binding, an application can either select a specific address, or bind to INADDR_ANY. Furthermore, on some systems other socket options (e.g., SO_BINDTODEVICE) can be used to bind to a specific

interface. If an application uses a specific address or binds to a specific interface, then MPTCP MUST respect this and not interfere in the application's choices. The binding to a specific address or interface implies that the application is not aware of MPTCP and will disable the use of MPTCP on this connection. An application that wishes to bind to a specific set of addresses with MPTCP must use multipath-aware calls to achieve this (as described in [Section 5.3.3](#)).

If an application binds to `INADDR_ANY`, it is assumed that the application does not care which addresses to use locally. In this case, a local policy MAY allow MPTCP to automatically set up multiple subflows on such a connection.

The basic sockets API of MPTCP-aware applications allows to express further preferences in an MPTCP-compatible way (e.g. bind to a subset of interfaces only).

[4.2.2](#). Querying of Addresses by Applications

Applications can use the `getpeername()` or `getsockname()` functions in order to retrieve the IP address of the peer or of the local socket. These functions can be used for various purposes, including security mechanisms, geo-location, or interface checks. The socket API was designed with an assumption that a socket is using just one address, and since this address is visible to the application, the application may assume that the information provided by the functions is the same during the lifetime of a connection. However, in MPTCP, unlike in TCP, there is a one-to-many mapping of a connection to subflows, and subflows can be added and removed while the connections continues to exist. Therefore, MPTCP cannot expose addresses by `getpeername()` or `getsockname()` that are both valid and constant during the connection's lifetime.

This problem is addressed as follows: If used by a legacy application, the MPTCP stack MUST always return the addresses of the first subflow of an MPTCP connection, in all circumstances, even if that particular subflow is no longer in use.

As this address may not be valid any more if the first subflow is closed, the MPTCP stack MAY close the whole MPTCP connection if the first subflow is closed (i.e. fate sharing between the initial subflow and the MPTCP connection as a whole). Whether to close the whole MPTCP connection by default SHOULD be controlled by a local policy. Further experiments are needed to investigate its implications.

The functions `getpeername()` and `getsockname()` SHOULD also always

return the addresses of the first subflow if the socket is used by an MPTCP-aware application, in order to be consistent with MPTCP-unaware applications, and, e. g., also with SCTP. Instead of `getpeername()` or `getsockname()`, MPTCP-aware applications can use new API calls, documented later, in order to retrieve the full list of address pairs for the subflows in use.

4.3. Socket Option Issues

4.3.1. General Guideline

The existing sockets API includes options that modify the behavior of sockets and their underlying communications protocols. Various socket options exist on socket, TCP, and IP level. The value of an option can usually be set by the `setsockopt()` system function. The `getsockopt()` function gets information. In general, the existing sockets interface functions cannot configure each MPTCP subflow individually. In order to be backward compatible, existing APIs therefore SHOULD apply to all subflows within one connection, as far as possible.

4.3.2. Disabling of the Nagle Algorithm

One commonly used TCP socket option (`TCP_NODELAY`) disables the Nagle algorithm as described in [2]. This option is also specified in the Posix standard [8]. Applications can use this option in combination with MPTCP exactly in the same way. It then SHOULD disable the Nagle algorithm for the MPTCP connection, i.e., all subflows.

In addition, the MPTCP protocol instance MAY use a different path scheduler algorithm if `TCP_NODELAY` is present. For instance, it could use an algorithm that is optimized for latency-sensitive traffic. Specific algorithms are outside the scope of this document.

4.3.3. Buffer Sizing

Applications can explicitly configure send and receive buffer sizes by the sockets API (`SO_SNDBUF`, `SO_RCVBUF`). These socket options can also be used in combination with MPTCP and then affect the buffer size of the MPTCP connection. However, when defining buffer sizes, application programmers should take into account that the transport over several subflows requires a certain amount of buffer for resequencing in the receiver. MPTCP may also require more storage space in the sender, in particular, if retransmissions are sent over more than one path. In addition, very small send buffers may prevent MPTCP from efficiently scheduling data over different subflows. Therefore, it does not make sense to use MPTCP in combination with small send or receive buffers.

An MPTCP implementation MAY set a lower bound for send and receive buffers and treat a small buffer size request as an implicit request not to use MPTCP.

4.3.4. Other Socket Options

Some network stacks also provide other implementation-specific socket options or interfaces that affect TCP's behavior. If a network stack supports MPTCP, it must be ensured that these options do not interfere.

4.4. Default Enabling of MPTCP

It is up to a local policy at the end system whether a network stack should automatically enable MPTCP for sockets even if there is no explicit sign of MPTCP awareness of the corresponding application. Such a choice may be under the control of the user through system preferences.

The enabling of MPTCP, either by application or by system defaults, does not necessarily mean that MPTCP will always be used. Both endpoints must support MPTCP, and there must be multiple addresses at at least one endpoint, for MPTCP to be used. Even if those requirements are met, however, MPTCP may not be immediately used on a connection. It may make sense for multiple paths to be brought into operation only after a given period of time, or if the connection is saturated.

4.5. Summary of Advices to Application Developers

- o Using the default MPTCP configuration: Like TCP, MPTCP is designed to be efficient and robust in the default configuration. Application developers should not explicitly configure TCP (or MPTCP) features unless this is really needed.
- o Socket buffer dimensioning: Multipath transport requires larger buffers in the receiver for resequencing, as already explained. Applications should use reasonable buffer sizes (such as the operating system default values) in order to fully benefit from MPTCP. A full discussion of buffer sizing issues is given in [5].
- o Facilitating stack-internal heuristics: The path management and data scheduling by MPTCP is realized by stack-internal algorithms that may implicitly try to self-optimize their behavior according to assumed application needs. For instance, an MPTCP implementation may use heuristics to determine whether an application requires delay-sensitive or bulk data transport, using for instance port numbers, the TCP_NODELAY socket options, or the

application's read/write patterns as input parameters. An application developer can facilitate the operation of such heuristics by avoiding atypical interface use cases. For instance, for long bulk data transfers, it does neither make sense to enable the TCP_NODELAY socket option, nor is it reasonable to use many small subsequent socket "send()" calls with small amounts of data only.

5. Basic API for MPTCP-aware Applications

5.1. Design Considerations

While applications can use MPTCP with the unmodified sockets API, multipath transport results in many degrees of freedom. MPTCP manages the data transport over different subflows automatically. By default, this is transparent to the application, but an application could use an additional API to interface with the MPTCP layer and to control important aspects of the MPTCP implementation's behaviour.

This document describes a basic MPTCP API. The API contains a minimum set of functions that provide an equivalent level of control and information as exists for regular TCP. It maintains backward compatibility with legacy applications.

An advanced MPTCP API is outside the scope of this document. The basic API does not allow a sender or a receiver to express preferences about the management of paths or the scheduling of data, even if this can have a significant performance impact and if an MPTCP implementation could benefit from additional guidance by applications. A list of potential further API extensions is provided in the appendix. The specification of such an advanced API is for further study and may partly be implementation-specific.

MPTCP mainly affects the sending of data. Therefore, the basic API only affects the sender side of a data transfer. A receiver may also have preferences about data transfer choices, and it may have performance requirements, too. A receiver may also have preferences about data transfer choices, and it may have performance requirements, too. Yet, the configuration of such preferences is outside of the scope of the basic API.

5.2. Requirements on the Basic MPTCP API

Because of the importance of the sockets interface there are several fundamental design objectives for the basic interface between MPTCP and applications:

- o Consistency with existing sockets APIs must be maintained as far as possible. In order to support the large base of applications using the original API, a legacy application must be able to continue to use standard socket interface functions when run on a system supporting MPTCP. Also, MPTCP-aware applications should be able to access the socket without any major changes.
- o Sockets API extensions must be minimized and independent of an implementation.
- o The interface should both handle IPv4 and IPv6.

The following is a list of the core requirements for the basic API:

- REQ1: Turn on/off MPTCP: An application should be able to request to turn on or turn off the usage of MPTCP. This means that an application should be able to explicitly request the use of MPTCP if this is possible. Applications should also be able to request not to enable MPTCP and to use regular TCP transport instead. This can be implicit in many cases, since MPTCP must be disabled by the use of binding to a specific address. MPTCP may also be enabled if an application uses a dedicated multipath address family (such as AF_MULTIPATH, [9]).
- REQ2: An application should be able to restrict MPTCP to binding to a given set of addresses.
- REQ3: An application should be able to obtain information on the addresses used by the MPTCP subflows.
- REQ4: An application should be able to extract a unique identifier for the connection (per endpoint).

The first requirement is the most important one, since some applications could benefit a lot from MPTCP, but there are also cases in which it hardly makes sense. The existing sockets API provides similar mechanisms to enable or disable advanced TCP features. The second requirement corresponds to the binding of addresses with the `bind()` socket call, or, e.g., explicit device bindings with a `SO_BINDTODEVICE` option. The third requirement ensures that there is an equivalent to `getpeername()` or `getsockname()` that is able to deal with more than one subflow. Finally, it should be possible for the application to retrieve a unique connection identifier (local to the endpoint on which it is running) for the MPTCP connection. This is equivalent to using the (address, port) pair for a connection identifier in single-path TCP, which is no longer static in MPTCP.

An application can continue to use `getpeername()` or `getsockname()` in addition to the basic MPTCP API. In that case, both functions return the corresponding addresses of the first subflow, as already explained.

5.3. Sockets Interface Extensions by the Basic MPTCP API

5.3.1. Overview

The abstract, basic MPTCP API consists of a set of new values that are associated with an MPTCP socket. Such values may be used for changing properties of an MPTCP connection, or retrieving information. These values could be accessed by new symbols on existing calls such as `setsockopt()` and `getsockopt()`, or could be implemented as entirely new function calls. This implementation decision is out of scope for this document. The following list presents symbolic names for these MPTCP socket settings.

- o `TCP_MULTIPATH_ENABLE`: Enable/disable MPTCP
- o `TCP_MULTIPATH_ADD`: Bind MPTCP to a set of given local addresses, or add a new local address to an existing MPTCP connection
- o `TCP_MUTLIPATH_REMOVE`: Remove a local address from a MPTCP connection
- o `TCP_MULTIPATH_SUBFLOWS`: Get the pairs of addresses currently used by the MPTCP subflows
- o `TCP_MULTIPATH_CONNID`: Get the local connection identifier for this MPTCP connection

Table Table 1 shows a list of the abstract socket operations for the basic configuration of MPTCP. The first column gives the symbolic name of the operation. The second and third columns indicate whether the operation provides values to be read ("Get") or takes values to configure ("Set"). The fourth column lists the type of data associated with this operation.

Name	Get	Set	Data type
TCP_MULTIPATH_ENABLE	o	o	boolean
TCP_MULTIPATH_ADD		o	list of addresses
TCP_MULTIPATH_REMOVE		o	list of addresses
TCP_MULTIPATH_SUBFLOWS	o		list of pairs of addresses
TCP_MULTIPATH_CONNID	o		32-bit integer

Table 1: MPTCP Socket Operations

There are restrictions when these new socket operations can be used:

- o TCP_MULTIPATH_ENABLE: This value SHOULD only be set before the establishment of a TCP connection. Its value SHOULD only be read after the establishment of a connection.
- o TCP_MULTIPATH_ADD: This operation can be both applied before connection setup or during a connection. If used before, it controls the local addresses that an MPTCP connection can use. In the latter case, it allows MPTCP to use an additional local address, if there has been a restriction before connection setup.
- o TCP_MULTIPATH_REMOVE: This operation only has meaning after connection setup.
- o TCP_MULTIPATH_SUBFLOWS: This value is read-only and SHOULD only be used after connection setup.
- o TCP_MULTIPATH_CONNID: This value is read-only and SHOULD only be used after connection setup.

5.3.2. Enabling and Disabling of MPTCP

An application can explicitly indicate multipath capability by setting TCP_MULTIPATH_ENABLE to a value larger than 0. In this case, the MPTCP implementation SHOULD try to negotiate MPTCP for that connection. Note that multipath transport will not necessarily be enabled, as it requires multiple addresses and support in the other end-system and potentially also on middleboxes.

An application can disable MPTCP setting TCP_MULTIPATH_ENABLE to a value of 0. In that case, MPTCP MUST NOT be used on that connection.

After connection establishment, an application can get the value of TCP_MULTIPATH_ENABLE. A value of 0 then means lack of MPTCP support. Any value equal to or larger than 1 means that MPTCP is supported.

As alternative to setting an explicit value, an application could also use a new, separate address family called `AF_MULTIPATH` [9]. This separate address family can be used to exchange multiple addresses between an application and the standard sockets API, and additionally acts as an explicit indication that an application is MPTCP-aware, i.e., that it can deal with the semantic changes of the sockets API, in particular concerning `getpeername()` and `getsockname()`. The usage of `AF_MULTIPATH` is also more flexible with respect to multipath transport, either IPv4 or IPv6, or both in parallel [9].

5.3.3. Binding MPTCP to Specified Addresses

Before connection establishment, an application can use `TCP_MULTIPATH_ADD` socket option to indicate a set of local IP addresses that MPTCP may bind to. By extension, this operation will also control the list of addresses that can be advertised to the peer via MPTCP signalling.

This operation can also be used to modify the address list in use during the lifetime of an MPTCP connection. In this case, it is used to indicate a set of additional local addresses that the MPTCP connection can make use of, and which can be signalled to the peer. It should be noted that this signal is only a hint, and an MPTCP implementation MAY only use a subset of the addresses.

The `TCP_MULTIPATH_REMOVE` operation can be used to remove a (set of) local addresses from an MPTCP connection. MPTCP MUST close any corresponding subflows (i.e. those using the local address that is no longer present), and signal the removal of the address to the peer. If alternative paths are available using the supplied address list but MPTCP is not currently using them, an MPTCP implementation SHOULD establish alternative subflows before undertaking the address removal.

It should be remembered that these operations SHOULD support both IPv4 and IPv6 addresses, potentially in the same call.

5.3.4. Querying the MPTCP Subflow Addresses

An application can get a list of the addresses used by the currently established subflows by means of the read-only `TCP_MULTIPATH_SUBFLOWS` operation. The return value is a list of pairs of IP addresses. In one pair, the first address refers to the local IP address, and the second one to the remote IP address used by the subflow. The list MUST only include established subflows. Each address pair MUST be either a pair of IPv4 addresses or a pair of IPv6 addresses.

5.3.5. Getting a Unique Connection Identifier

An application that wants a unique identifier for the connection, analogous to an (address, port) pair in regular TCP, can query the `TCP_MULTIPATH_CONNID` value to get a local connection identifier for the MPTCP connection.

This is a 32-bit number, and SHOULD be the same as the local connection identifier sent in the MPTCP handshake.

6. Other Compatibility Issues

6.1. Usage of the Sctp Socket API

For dealing with multi-homing, several socket API extensions have been defined for Sctp [[14](#)]. As MPTCP realizes multipath transport from and to multi-homed endsystems, some of these interface function calls are actually applicable to MPTCP in a similar way.

The following functions that are defined for Sctp have similar functionality to the MPTCP API extensions defined earlier:

- o `sctp_bindx()`
- o `sctp_connectx()`
- o `sctp_getladdrs()`
- o `sctp_getpaddrs()`

The syntax and semantics of these functions are described in [[14](#)].

API developers MAY wish to integrate Sctp and MPTCP calls to provide a consistent interface to the application. Yet, it must be emphasized that the transport service provided by MPTCP is different to Sctp, and this is why not all Sctp API functions can be mapped directly to MPTCP. Furthermore, a network stack implementing MPTCP does not necessarily support Sctp and its specific socket interface extensions. This is why the basic API of MPTCP defines additional socket options only, which are a backward compatible extension of TCP's application interface.

6.2. Incompatibilities with other Multihoming Solutions

The use of MPTCP can interact with various related sockets API extensions. The use of a multihoming shim layer conflicts with multipath transport such as MPTCP or Sctp [[12](#)]. Care should be taken for the usage not to confuse with the overlapping features of other

APIs:

- o SHIM API [[12](#)]: This API specifies sockets API extensions for the multihoming shim layer.
- o HIP API [[13](#)]: The Host Identity Protocol (HIP) also results in a new API.
- o API for Mobile IPv6 [[11](#)]: For Mobile IPv6, a significantly extended socket API exists as well.

In order to avoid any conflict, multiaddressed MPTCP SHOULD NOT be enabled if a network stack uses SHIM6, HIP, or Mobile IPv6. Furthermore, applications should not try to use both the MPTCP API and another multihoming or mobility layer API.

It is possible, however, that some of the MPTCP functionality, such as congestion control, could be used in a SHIM6 or HIP environment. Such operation is outside the scope of this document.

[6.3.](#) Interactions with DNS

In multihomed or multiaddressed environments, there are various issues that are not specific to MPTCP, but have to be considered, too. These problems are summarized in [[15](#)].

Specifically, there can be interactions with DNS. Whilst it is expected that an application will iterate over the list of addresses returned from a call such as `getaddrinfo()`, MPTCP itself MUST NOT make any assumptions about multiple A or AAAA records from the same DNS query referring to the same host, as it is possible that multiple addresses refer to multiple servers for load balancing purposes.

TODO: Elaborate on DNS

[7.](#) Security Considerations

Will be added in a later version of this document.

[8.](#) IANA Considerations

No IANA considerations.

[9.](#) Conclusion

This document discusses MPTCP's application implications and specifies a basic MPTCP API. For legacy applications, it is ensured that the existing sockets API continues to work. MPTCP-aware

applications can use the basic MPTCP API that provides some control over the transport layer equivalent to regular TCP. A more fine-granular interaction between applications and MPTCP requires an advanced MPTCP API, which is not specified in this document.

10. Acknowledgments

Authors sincerely thank to the following people for their helpful comments to the document: Costin Raiciu, Philip Eardley

Michael Scharf is supported by the German-Lab project (<http://www.german-lab.de/>) funded by the German Federal Ministry of Education and Research (BMBF). Alan Ford is supported by Trilogy (<http://www.trilogy-project.org/>), a research project (ICT-216372) partially funded by the European Community under its Seventh Framework Program. The views expressed here are those of the author(s) only. The European Commission is not liable for any use that may be made of the information in this document.

11. References

11.1. Normative References

- [1] Postel, J., "Transmission Control Protocol", STD 7, [RFC 793](#), September 1981.
- [2] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, [RFC 1122](#), October 1989.
- [3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [4] Ford, A., Raiciu, C., Handley, M., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development", [draft-ietf-mptcp-architecture-02](#) (work in progress), October 2010.
- [5] Ford, A., Raiciu, C., and M. Handley, "TCP Extensions for Multipath Operation with Multiple Addresses", [draft-ietf-mptcp-multiaddressed-02](#) (work in progress), October 2010.
- [6] Bagnulo, M., "Threat Analysis for Multi-addressed/Multi-path TCP", [draft-ietf-mptcp-threat-03](#) (work in progress), October 2010.
- [7] Raiciu, C., Handley, M., and D. Wischik, "Coupled Multipath-Aware Congestion Control", [draft-ietf-mptcp-congestion-00](#) (work

in progress), July 2010.

- [8] "IEEE Std. 1003.1-2008 Standard for Information Technology -- Portable Operating System Interface (POSIX). Open Group Technical Standard: Base Specifications, Issue 7, 2008."

11.2. Informative References

- [9] Sarolahti, P., "Multi-address Interface in the Socket API", [draft-sarolahti-mptcp-af-multipath-01](#) (work in progress), March 2010.
- [10] Stevens, W., Thomas, M., Nordmark, E., and T. Jinmei, "Advanced Sockets Application Program Interface (API) for IPv6", [RFC 3542](#), May 2003.
- [11] Chakrabarti, S. and E. Nordmark, "Extension to Sockets API for Mobile IPv6", [RFC 4584](#), July 2006.
- [12] Komu, M., Bagnulo, M., Slavov, K., and S. Sugimoto, "Socket Application Program Interface (API) for Multihoming Shim", [draft-ietf-shim6-multihome-shim-api-13](#) (work in progress), February 2010.
- [13] Komu, M. and T. Henderson, "Basic Socket Interface Extensions for Host Identity Protocol (HIP)", [draft-ietf-hip-native-api-12](#) (work in progress), January 2010.
- [14] Stewart, R., Poon, K., Tuexen, M., Yasevich, V., and P. Lei, "Sockets API Extensions for Stream Control Transmission Protocol (SCTP)", [draft-ietf-tsvwg-sctpsocket-23](#) (work in progress), July 2010.
- [15] Blanchet, M. and P. Seite, "Multiple Interfaces Problem Statement", [draft-ietf-mif-problem-statement-04](#) (work in progress), May 2010.
- [16] Wasserman, M. and P. Seite, "Current Practices for Multiple Interface Hosts", [draft-ietf-mif-current-practices-01](#) (work in progress), June 2010.

Appendix A. Requirements on a Future Advanced MPTCP API

A.1. Design Considerations

Multipath transport results in many degrees of freedom. The basic MPTCP API only defines a minimum set of the API extensions for the interface between the MPTCP layer and applications, which does not

offer much control of the MPTCP implementation's behaviour. A future, advanced API could address further features of MPTCP and provide more control.

Applications that use TCP may have different requirements on the transport layer. While developers have become used to the characteristics of regular TCP, new opportunities created by MPTCP could allow the service provided to be optimised further. An advanced API could enable MPTCP-aware applications to specify preferences and control certain aspects of the behavior, in addition to the simple control provided by the basic interface. An advanced API could also address aspects that are completely out-of-scope of the basic API, for example, the question whether a receiving application could influence the sending policy.

Furthermore, an advanced MPTCP API could be part of a new overall interface between the network stack and applications that addresses other issues as well, such as the split between identifiers and locators. An API that does not use IP addresses (but, instead e.g. a `connectbyname()` function) would be useful for numerous purposes, independent of MPTCP.

This appendix documents a list of potential usage scenarios and requirements for the advanced API. The specification and implementation of a corresponding API is outside the scope of this document.

[A.2.](#) MPTCP Usage Scenarios and Application Requirements

There are different MPTCP usage scenarios. An application that wishes to transmit bulk data will want MPTCP to provide a high throughput service immediately, through creating and maximising utilisation of all available subflows. This is the default MPTCP use case.

But at the other extreme, there are applications that are highly interactive, but require only a small amount of throughput, and these are optimally served by low latency and jitter stability. In such a situation, it would be preferable for the traffic to use only the lowest latency subflow (assuming it has sufficient capacity), maybe with one or two additional subflows for resilience and recovery purposes. The key challenge for such a strategy is that the delay on a path may fluctuate significantly and that just always selecting the path with the smallest delay might result in instability.

The choice between bulk data transport and latency-sensitive transport affects the scheduler in terms of whether traffic should be, by default, sent on one subflow or across several ones. Even if

the total bandwidth required is less than that available on an individual path, it is desirable to spread this load to reduce stress on potential bottlenecks, and this is why this method should be the default for bulk data transport. However, that may not be optimal for applications that require latency/jitter stability.

In the case of the latter option, a further question arises: Should additional subflows be used whenever the primary subflow is overloaded, or only when the primary path fails (hot-standby)? In other words, is latency stability or bandwidth more important to the application? This results in two different options: Firstly, there is the single path which can overflow into an additional subflow; and secondly there is single-path with hot-standby, whereby an application may want an alternative backup subflow in order to improve resilience. In case that data delivery on the first subflow fails, the data transport could immediately be continued on the second subflow, which is idle otherwise.

Yet another complication is introduced with the potential that MPTCP introduces for changes in available bandwidth as the number of available subflows changes. Such jitter in bandwidth may prove confusing for some applications such as video or audio streaming that dynamically adapt codecs based on available bandwidth. Such applications may prefer MPTCP to attempt to provide a consistent bandwidth as far as is possible, and avoid maximising the use of all subflows.

A further, mostly orthogonal question is whether data should be duplicated over the different subflows, in particular if there is spare capacity. This could improve both the timeliness and reliability of data delivery.

In summary, there are at least three possible performance objectives for multipath transport (not necessarily disjoint):

1. High bandwidth
2. Low latency and jitter stability
3. High reliability

In an advanced API, applications could provide high-level guidance to the MPTCP implementation concerning these performance requirements, for instance, which is considered to be the most important one. The MPTCP stack would then use internal mechanisms to fulfill this abstract indication of a desired service, as far as possible. This would both affect the assignment of data (including retransmissions) to existing subflows (e.g., 'use all in parallel', 'use as overflow',

'hot standby', 'duplicate traffic') as well as the decisions when to set up additional subflows to which addresses. In both cases different policies can exist, which can be expected to be implementation-specific.

Therefore, an advanced API could provide a mechanism how applications can specify their high-level requirements in an implementation-independent way. One possibility would be to select one "application profile" out of a number of choices that characterize typical applications. Yet, as applications today do not have to inform TCP about their communication requirements, it requires further studies whether such an approach would be realistic.

Of course, independent of an advanced API, such functionality could also partly be achieved by MPTCP-internal heuristics that infer some application preferences e.g. from existing socket options, such as TCP_NODELAY. Whether this would be reliable, and indeed appropriate, is for further study, too.

A.3. Potential Requirements on an Advanced MPTCP API

The following is a list of potential requirements for an advanced MPTCP API beyond the features of the basic API. It is included here for information only:

- REQ5: An application should be able to establish MPTCP connections without using IP addresses as locators.
- REQ6: An application should be able obtain usage information and statistics about all subflows (e.g., ratio of traffic sent via this subflow).
- REQ7: An application should be able to request a change in the number of subflows in use, thus triggering removal or addition of subflows. An even finer control granularity would be a request for the establishment of a new subflow to a provided destination, or a request for the termination of a specified, existing subflow.
- REQ8: An application should be able to inform the MPTCP implementation about its high-level performance requirements, e.g., in form of a profile.
- REQ9: An application should be able to control the automatic establishment/termination of subflows. This would imply a selection among different heuristics of the path manager, e.g., 'try as soon as possible', 'wait until there is a bunch of data', etc.

REQ10: An application should be able to set preferred subflows or subflow usage policies. This would result in a selection among different configurations of the multipath scheduler.

REQ11: An application should be able to control the level of redundancy by telling whether segments should be sent on more than one path in parallel.

An advanced API fulfilling these requirements would allow application developers to more specifically configure MPTCP. It could avoid suboptimal decisions of internal, implicit heuristics. However, it is unclear whether all of these requirements would have a significant benefit to applications, since they are going above and beyond what the existing API to regular TCP provides.

[Appendix B](#). Change History of the Document

Changes compared to version 03:

- o Removal of explicit references to "socket options" and `getsockopt/setsockopt`.
- o Change of `TCP_MULTIPATH_BIND` to `TCP_MULTIPATH_ADD` and `TCP_MULTIPATH_REMOVE`.
- o Mention of stability of bandwidth as another potential QoS parameter for the advanced API.
- o Address comments received from Philip Eardley: Explanation of the API terminology, more explicit statement concerning applications that bind to a specific address, and some smaller editorial fixes

Changes compared to version 02:

- o Definition of the behavior of `getpeername()` and `getsockname()` when being called by an MPTCP-aware application.
- o Discussion of the possibility that an MPTCP implementation could support the SCTP API, as far as it is applicable to MPTCP.
- o Various editorial fixes.

Changes compared to version 01:

- o Second half of the document completely restructured
- o Separation between a basic API and an advanced API: The focus of the document is the basic API only; all text concerning a

potential extended API is moved to the appendix

- o Several clarifications, e. g., concerning buffer sizeing and the use of different scheduling strategies triggered by TCP_NODELAY
- o Additional references

Changes compared to version 00:

- o Distinction between legacy and MPTCP-aware applications
- o Guidance concerning default enabling, reaction to the shutdown of the first subflow, etc.
- o Reference to a potential use of AF_MULTIPATH
- o Additional references to related work

Authors' Addresses

Michael Scharf
Alcatel-Lucent Bell Labs
Lorenzstrasse 10
70435 Stuttgart
Germany

E-Mail: michael.scharf@alcatel-lucent.com

Alan Ford
Roke Manor Research
Old Salisbury Lane
Romsey, Hampshire S051 0ZN
UK

Phone: +44 1794 833 465
E-Mail: alan.ford@roke.co.uk

