

MPTCP Working Group
Internet-Draft
Intended status: Informational
Expires: May 23, 2020

S. Barre
G. Detal
Tessares
O. Bonaventure
UCLouvain and Tessares
C. Paasch
Apple
November 20, 2019

TFO support for Multipath TCP
draft-barre-mptcp-tfo-06

Abstract

TCP Fast Open (TFO) is a TCP extension that allows sending data in the SYN, instead of waiting until the TCP connection is established. This document describes what parts of Multipath TCP must be adapted to support it, and how TFO and MPTCP can operate together.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on May 23, 2020.

Copyright Notice

Copyright (c) 2019 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 (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must

include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. TFO cookie request with MPTCP	3
3. Data sequence mapping under TFO	4
4. Early context creation in server	4
5. Using TFO to avoid useless MPTCP negotiations	5
6. Using TFO with MP_JOIN	6
7. Connection establishment examples	6
8. Middlebox interactions	8
9. Security considerations	9
10. Conclusion	10
11. Acknowledgements	10
12. Informative References	10
Appendix A. Implementation status	10
Authors' Addresses	11

1. Introduction

TCP Fast Open, described in [RFC7413], has been introduced with the objective of gaining one RTT before transmitting data. This is considered a valuable gain as very short connections are very common, especially for HTTP request/response schemes. MPTCP, on the other hand, has been defined in [I-D.ietf-mptcp-rfc6824bis] to add multipath support to TCP, where a TCP flow is divided in several TCP subflows. Given that MPTCP can be applied transparently to any TCP socket, without the application knowing, it should be able to support TCP fast open when the application asks for it.

When doing that, one important thing to examine is the option length consumed in segments that would carry both a TFO and an MPTCP option. The handling of MPTCP data sequence mappings must also be updated to take into account the data that is sent together with the SYN or the SYN+ACK. A third issue to handle is the state creation in the server: TFO allows the server to create TCP state as soon as a SYN is received. With MPTCP, even more state is created, and it may be useful to avoid this in a situation where MPTCP does not work but TFO does.

The rest of this document is organized as follows:

Section 2 describes the TFO cookie request, in the case of a Multipath TCP flow. Section 3 proposes a way to map SYN data to the data sequence number space, while taking middleboxes into account.

In Section 4, it is explained that the MP_CAPABLE option is no longer always necessary in the third ack of the three-way handshake. Section 5 presents two ways to avoid useless MPTCP context creations in the server, one for client implementations, the other for server implementations, as a TFO extension. Section 6 takes the MP_JOIN case into consideration. Finally, we describe middlebox interactions in Section 8, and security considerations in Section 9.

2. TFO cookie request with MPTCP

When a TFO client first connects to a server, it cannot immediately include data in the SYN for security reasons [RFC7413]. Instead, it requests a cookie that will be used in subsequent connections. This is done with the TCP cookie request/response options, of resp. 2 bytes and 6-18 bytes (depending on the chosen cookie length).

TFO and MPTCP can be combined provided that the total length of their options does not exceed the maximum 40 bytes possible in TCP:

- o In the SYN: MPTCP uses a 4-bytes long MP_CAPABLE option. The MPTCP and TFO options sum up to 6 bytes. [I-D.ietf-mptcp-rfc6824bis] mentions in Appendix A that SYN packet options typically sum up to 19 bytes, or 24 bytes where implementations pad each option up to a word boundary. Even in the worst case, this fits the maximum option space.
- o In the SYN+ACK: MPTCP uses a 12-bytes long MP_CAPABLE option, but now TFO can be as long as 18 bytes. Since the maximum option length may be exceeded, it is up to the server to solve this by using a shorter cookie or pad the whole option block instead of each option separately. Alternatively, the server may decide to fallback to MPTCP-only (by not giving a cookie at all), or to TFO-only. As an example, if we consider that 19 bytes are used for classical TCP options, the maximum possible cookie length would be of 7 bytes. The consequence of this, from a security viewpoint, is explored in Section 9. Note that the same limitation applies to subsequent connections, for the SYN packet (because the client then echoes back the cookie to the server). Finally, if the security impact of reducing the cookie size is not deemed acceptable, the server can reduce the amount of other TCP-options by omitting the TCP timestamps. Indeed, as outlined in [I-D.ietf-mptcp-rfc6824bis] in Appendix A, an MPTCP connection could also avoid the use of TCP timestamps thanks to MPTCP's use of 64-bit sequence numbers which already provides protection against wrapped sequence numbers.
- o In the third ACK: Nothing special compared to MPTCP, since no TFO option is used there.

Once the cookie has been successfully exchanged, the rest of the connection is just regular MPTCP. The rest of this document assumes that the cookie request has been exchanged, and that data can be included in the SYN.

3. Data sequence mapping under TFO

MPTCP [I-D.ietf-mptcp-rfc6824bis] uses, in the TCP establishment phase, a key exchange that is used to generate the Initial Data Sequence Numbers (IDSNs). More precisely, [I-D.ietf-mptcp-rfc6824bis] states in section 3.1 that "The SYN with MP_CAPABLE occupies the first octet of data sequence space, although this does not need to be acknowledged at the connection level until the first data is sent". With TFO, one way to handle the data sent together with the SYN would be to consider an implicit DSS mapping that covers that SYN segment (since there is not enough space in the SYN to include a DSS option). The problem with that approach is that if a middlebox modifies the TFO data, this will not be noticed by MPTCP because of the absence of a DSS-checksum. For example, a TCP (but not MPTCP)-aware middlebox could insert bytes at the beginning of the stream and adapt the TCP checksum and sequence numbers accordingly. With an implicit mapping, this would give to client and server a different view on the DSS-mapping, with no way to detect this inconsistency as the DSS checksum is not present. One way to solve this is to simply consider that the TFO data is not part of the Data Sequence Number space: the SYN with MP_CAPABLE still occupies the first octet of data sequence space, but then the first non-TFO data byte occupies the second octet. This guarantees that, if the use of DSS-checksum is negotiated, all data in the data sequence number space is checksummed. We also note that this does not entail a loss of functionality, because TFO-data is always sent when only one path is active.

4. Early context creation in server

In order to enable the server to receive and send data before the end of the three-way handshake, TFO allows creating state on the server as soon as the SYN is received if a valid cookie is provided. The MPTCP state should then also be created upon SYN reception (see exceptions for that in Section 5).

DISCUSSION: Doing that allows relaxing the MPTCP MP_CAPABLE exchange, in that the sender's and receiver's keys are no longer required in the third ack of the three-way handshake, because their role was precisely to compensate for the absence of server state until the end of the establishment. The consequence is that the MP_CAPABLE option can simply be removed from the third ack. However, an MPTCP option must still be present when concluding the three-way handshake, to

confirm to the server that its own MP_CAPABLE option (in the SYN+ACK) has been correctly received by the client. The DSS option can replace the MP_CAPABLE option, while simultaneously allowing the transmission of more data in the third ack. Moreover, providing a DSS option to the server early allows faster establishment of new subflows (see [I-D.ietf-mptcp-rfc6824bis], Section 3.1).

In order to decide whether it can send a third ack with DSS-only instead of MP_CAPABLE, a client must verify if the TFO data has been at least partially acknowledged. If the SYN+ACK only acknowledges the SYN, TFO may be not supported in the server, or the cookie may have been filtered by the network. There is no guarantee that the MPTCP state has been created, and the third ack should contain the MP_CAPABLE option, with the client and server keys.

5. Using TFO to avoid useless MPTCP negotiations

The TFO cookie, sent in a SYN, indicates that a previous connection has been successfully established, and that TCP state can safely be created. It does not however say anything about whether the MPTCP options are filtered or not in the network. It is thus possible that a server creates an MPTCP context upon SYN+TFO cookie reception, then actually needs to discard it after having discovered that the MPTCP options are filtered.

One way to solve this would be for the client to cache destinations that do support MPTCP. TFO allows sending data together with the SYN starting at the second connection. The first one is used to learn the cookie from the server. It could also be used to learn whether MPTCP can be used with the peer.

DISCUSSION: The other, compatible way to solve the problem is to extend TFO and cache the Multipath Capability in the cookie generated by the server. The server could modify its cookie computation, to include multipath capability information in the cookie. Then, upon SYN+TFO cookie reception, the server could easily determine if the initial TFO flow was a successful MPTCP connection or not. The problem with this approach is that the server does not know yet whether the flow is multipath-capable when sending the TFO cookie. It could then send a first pessimistic cookie, as `GetCookie(IP_Address, mp_capable=false)` (adapted from [RFC7413], Section 4.1.2). Then, when it is determined that the flow is Multipath Capable (third ack received with an MPTCP option), a new cookie=`GetCookie(IP_Address, mp_capable=true)` can be generated and sent in the FIN to ensure reliable delivery.

6. Using TFO with MP_JOIN

TFO must not be used when establishing joined subflows. Doing that would be in contradiction with [I-D.ietf-mptcp-rfc6824bis], that states in section 3.2 that "It is not permitted to send data while in the PRE_ESTABLISHED state". Using TFO with joined subflows would mean that data is sent even before getting to the PRE_ESTABLISHED state.

7. Connection establishment examples

In this section we show a few examples of possible TFO+MPTCP establishment scenarios. For representing segments, we use the Tcpdump syntax.

Before a client can send data together with the SYN, it must request a cookie to the server, as shown in Figure 1. This is done by simply combining the TFO and MPTCP options.

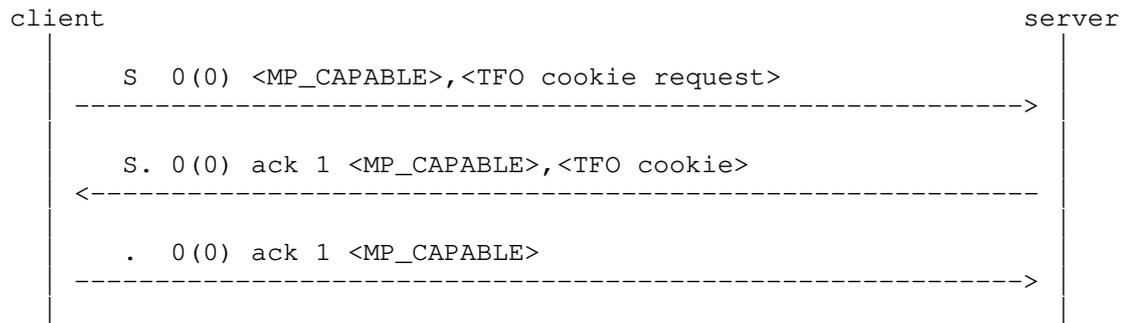


Figure 1: Cookie request

Once this is done, the received cookie can be used for TFO, as shown in Figure 2. The MP_CAPABLE is no longer required for the third ack, as explained in Section 4. Note that the last segment in the figure has a TCP sequence number of 21, while the DSS subflow sequence number is 1 (because the TFO data is not part of the data sequence number space, as explained in Section 3).

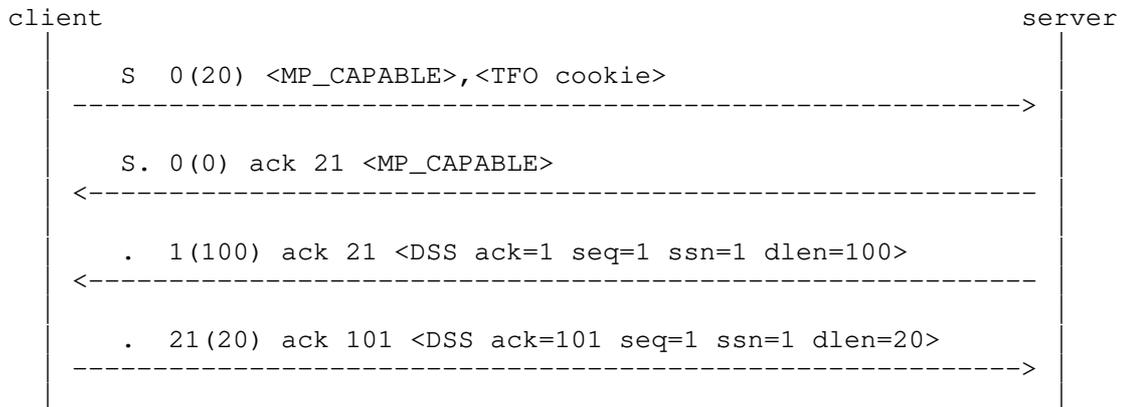


Figure 2: The server supports TFO

In Figure 3, the server does not support TFO. The client detects that no state is created in the server (as no data is acked), and now sends the MP_CAPABLE in the third ack, in order for the server to build its MPTCP context at then end of the establishment. Now, the tfo data, retransmitted, becomes part of the data sequence mapping because it is effectively sent (in fact re-sent) after the establishment.

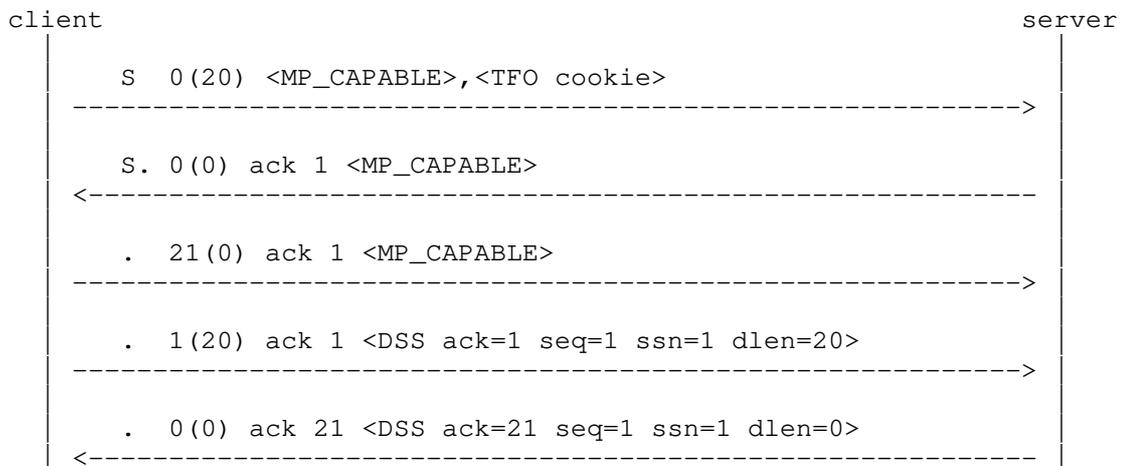


Figure 3: The server does not support TFO

It is also possible that the server acknowledges only part of the TFO data, as illustrated in Figure 4.

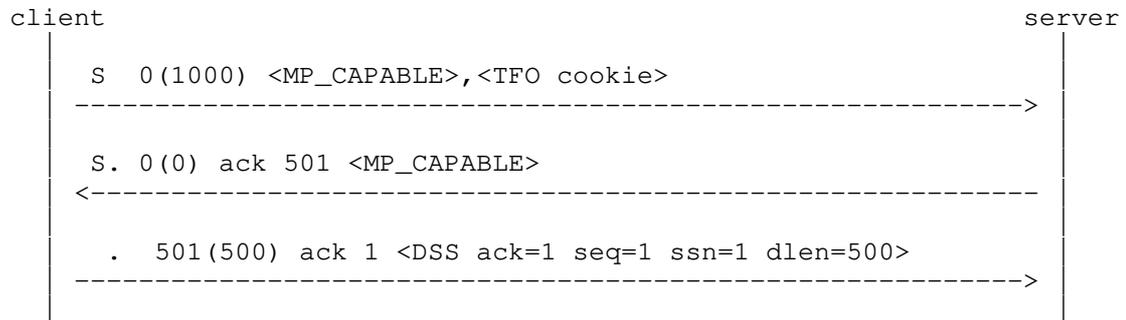


Figure 4: Partial data acknowledgement

8. Middlebox interactions

[I-D.ietf-mptcp-rfc6824bis], Section 6, describes middlebox interactions for Multipath TCP. This document does not define any new option compared to MPTCP or TFO. It defines a combination of them.

TFO also defines how an implementation should react when the TFO SYN is lost (fallback to regular TCP, [RFC7413] Section 4.2.1).

We propose to remove the MP_CAPABLE option from the third ack when TFO is used, based on the assumption that the context has been created already in the server upon SYN reception. Should the server actually not create this state, it would not be able to create its MPTCP state and would fallback to regular TCP. The state is not created in the server if it has no TFO support or the cookie is invalid, but in that case only the SYN is acknowledged, and the client does send the MP_CAPABLE option.

The other case where the server does not create MPTCP state is when the cookie includes a "mp_capable=false" information. In that case, regular TCP is used to take into account middleboxes that prevent correct MPTCP operation.

Even though this document presents mechanisms for collaboration between MPTCP and TFO, the filtering of one will not stop the other from working. For example, if a TFO option is dropped, MPTCP will fallback to sending MP_CAPABLE in third_ack, because no TFO data is acked. If the server stores MPTCP information in the cookie, this will be completely opaque to the network, and even to the client. Should that cookie be transformed or lost, it would not be accepted anymore by the server, which would fallback to regular MPTCP

communication, or regular TCP if MPTCP options are also filtered or modified.

The problem of middleboxes that alter the TFO data is solved by the fact that TFO data is not part of the Data Sequence Number space, as explained in Section 3.

9. Security considerations

Compared to using TFO or MPTCP alone, implementing the present combination could lead to more state created in the server, since MPTCP now creates state as soon as the first SYN is received. This is however not considered as a problem, for the following reasons:

- o The server will only create state when a valid TFO cookie is received. This guarantees that a successful TCP connection has been previously established with the same peer.
- o It remains possible that a useless MPTCP context is created upon SYN reception (due to TFO support but MPTCP options being filtered by the network). This is more an optimization issue than a security issue given the TFO cookie protection already present. Section 5 still proposes a solution to avoid creating MPTCP state in that case.
- o When under memory pressure, a server always has the option to refuse the client cookie. In that case, the session establishment will happen without data, and the client will send the MP_CAPABLE option in the third ack so that the server can create the MPTCP context at that time.

As mentioned in Section 2, it may be required to reduce the length of the cookie when MPTCP and TFO are used together. This can become a security issue when attackers and networks become fast enough for a brute force attack to be successful. An option to solve this would be to use TCP payload to store additional options, as suggested in [I-D.ietf-mptcp-rfc6824bis], Section 5. Another way would be to allow longer TCP options by using an "Extended Data Offset Option" [I-D.touch-tcpm-tcp-edo]. The problem with this is that the most problematic segment in the present case is the SYN (with long TFO cookie and MP_CAPABLE MPTCP option), for which it is more difficult to apply the Extended Data Offset Option ([I-D.touch-tcpm-tcp-edo], Section 7.7).

10. Conclusion

In this document, we have proposed minor extensions to MPTCP and TFO to allow them to operate together. In particular, we proposed excluding the TFO data from the data sequence number space. We explained that TFO allows to relax the MPTCP establishment in that the MP_CAPABLE option of the third ack can be removed in some cases. We also emphasized that such a combination augments the size of the TCP options, already quite large, although the combination is still possible with common TCP options and limited cookie length. We also proposed a way to cache multipath capability information in the client or in the TFO cookie. Finally, we examined potential middlebox interaction problems, or security problems that would arise from that combined operation.

11. Acknowledgements

This work was supported by the FP7-Trilogy2 project and by the Belgian Walloon Region under its FIRST Spin-Off Program (RICE project).

12. Informative References

[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-18 (work in progress), June 2019.

[I-D.touch-tcpm-tcp-edo]

Touch, J. and W. Eddy, "TCP Extended Data Offset Option", draft-touch-tcpm-tcp-edo-03 (work in progress), July 2014.

[MultipathTCP-Linux]

Paasch, C., Barre, S., and . et al, "Multipath TCP in the Linux kernel", n.d., <<http://www.multipath-tcp.org>>.

[RFC7413]

Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, DOI 10.17487/RFC7413, December 2014, <<https://www.rfc-editor.org/info/rfc7413>>.

Appendix A. Implementation status

In this section, we present the report of the implementation of this draft inside the Linux reference implementation of Multipath TCP [MultipathTCP-Linux]. The support of TFO in the MPTCP stack has been implemented on the 3.14 kernel (MPTCP v0.89).

The main design choices of this implementation are the following:

- o Minimize the modification to the current MPTCP and TFO stacks, i.e. let the TFO stack deal with sending data, receiving data inside the SYN.
- o Create the MPTCP state when receiving a SYN with a valid token on the server side as defined in Section 4.
- o Map the remaining data segments in the receive and send buffers to MPTCP data sequence numbers.

This latter point needs further explanation. First, in the current reference implementation of MPTCP, the MPTCP state is created upon reception of the SYN+ACK on the client-side. The implementation however did the MPTCP state allocation before processing the actual acknowledgement at the subflow level. This means that data (even acknowledged by the SYN+ACK) remains in the send buffer at the time of the allocation (which contained only the SYN in the case of regular MPTCP). We modified this behaviour to ensure that only unacknowledged data remains in the send buffer when allocating the state. Moreover, as the data was initially sent over the regular TCP flow, they had no MPTCP sequence numbers (the MPTCP state did not exist during the initial `sendto()` call). After the allocation of the MPTCP state, we modify these sequence numbers such that they are mapped starting at "IDSN + 1". This effectively gives the data sequence number "IDSN + 1" to the first byte following the establishment, since the acknowledged TFO data has been removed from the queues at this point. This data will then follow the same path as for data sent via a regular `write()` call.

As is the case for unacknowledged data on the client-side, the server-side can also have data in the receive buffer (the data sent in the SYN). We perform the same operation by mapping this data from TCP to MPTCP sequence numbers. TFO data is then mapped ahead of the IDSN, so as to ensure, again, that the first byte following the establishment has the data sequence number "IDSN + 1".

As of this writing, the implementation still generates a regular third acknowledgment with a `MP_CAPABLE` option (see Section 4) and it does not take benefit from the TFO cache to avoid useless MPTCP negotiation (see Section 5).

Authors' Addresses

Sebastien Barre
Tessares

Email: sebastien.barre@tessares.net

Gregory Detal
Tessares

Email: gregory.detal@tessares.net

Olivier Bonaventure
UCLouvain and Tessares

Email: Olivier.Bonaventure@tessares.net

Christoph Paasch
Apple

Email: cpaasch@apple.com

MPTCP Working Group
Internet-Draft
Updates: RFC6824 (if approved)
Intended status: Informational
Expires: January 7, 2016

O. Bonaventure
UCLouvain
July 06, 2015

Multipath TCP Address Advertisement
draft-bonaventure-mptcp-addr-00

Abstract

Multipath TCP [RFC6824] defines the ADD_ADDR option to allow a host to announce its addresses to the remote host. In this document we analyze some of the issues with the address advertisement technique defined [RFC6824] and propose some modifications to mitigate these problems. We also show that the reverse DNS could be an excellent alternative to advertise the stable addresses of a server.

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 January 7, 2016.

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

1. Introduction	2
2. Issues with ADD_ADDR	3
2.1. Usage of the Address Identifiers	4
2.2. Reliability of the ADD_ADDR Option	5
3. Learning the Addresses bound to a host through the DNS	5
4. Conclusion	7
5. Acknowledgements	7
6. Informative References	7
Author's Address	9

1. Introduction

Multipath TCP is an extension to TCP [RFC0793] that was specified in [RFC6824]. Multipath TCP was designed with multi-addressed hosts in mind [RFC6182]. A Multipath TCP connection is composed of subflows that are established between any of the addresses of the communicating hosts. [RFC6824] defines two options to manage the host addresses :

- o ADD_ADDR is used to announce one address bound to a host (possibly combined with a port number)
- o REMOVE_ADDR is used to indicate that an address previously attached to a host is not anymore attached to this host

To cope with Network Address Translation (NAT), the ADD_ADDR and REMOVE_ADDR options contain an address identifier encoded as an 8 bits integer.

When the initial subflow is created, it is assumed to be initiated from the address of the client whose identifier is 0 towards the address of the server whose identifier is also 0. Both the client and the server can use ADD_ADDR to advertise the other addresses that they use. When an additional subflow is created, the MP_JOIN option placed in the SYN (resp. SYN+ACK) contains the identifier of the address used to create (resp. accept) the subflow.

Experience with Multipath TCP shows that these two options allow to support multi-homed or dual-stack servers [TMA2015] and mobile devices [Cellnet12]. While the ADD_ADDR option has been supported in the Linux implementation of Multipath TCP, other implementors have chosen to not support it [I-D.eardley-mptcp-implementations-survey] while still supporting the REMOVE_ADDR option.

In this document, we first analyse in Section 2 several issues with the current ADD_ADDR option as defined in [RFC6824] and [I-D.ietf-mptcp-rfc6824bis]. Then in Section 3 we show how Multipath TCP could leverage the existing DNS to obtain information about the different addresses attached to a server.

2. Issues with ADD_ADDR

A first issue are the security risks if an attacker is able to send spoofed TCP segments that include the ADD_ADDR option. Multipath TCP [RFC6824] defines the ADD_ADDR option shown in Figure 1.

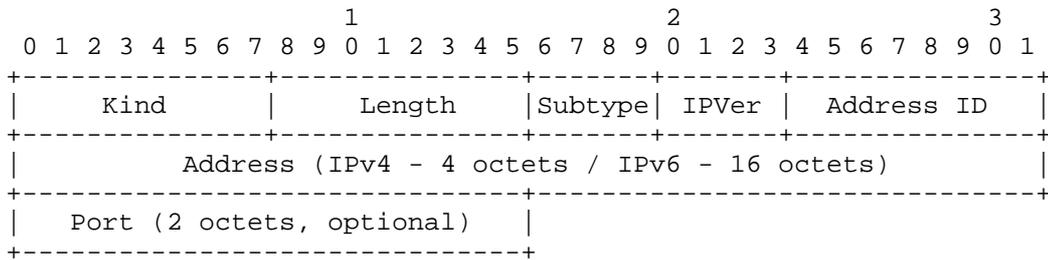


Figure 1: The ADD_ADDR option

From a security viewpoint, this option introduces a potential security risk if an attacker is able to send a spoofed ADD_ADDR option. [I-D.ietf-mptcp-rfc6824bis] proposes a new format for this option by placing a truncated HMAC inside the option to authenticate it. The format for this new option (ADD_ADDR2) is shown in Figure 2.

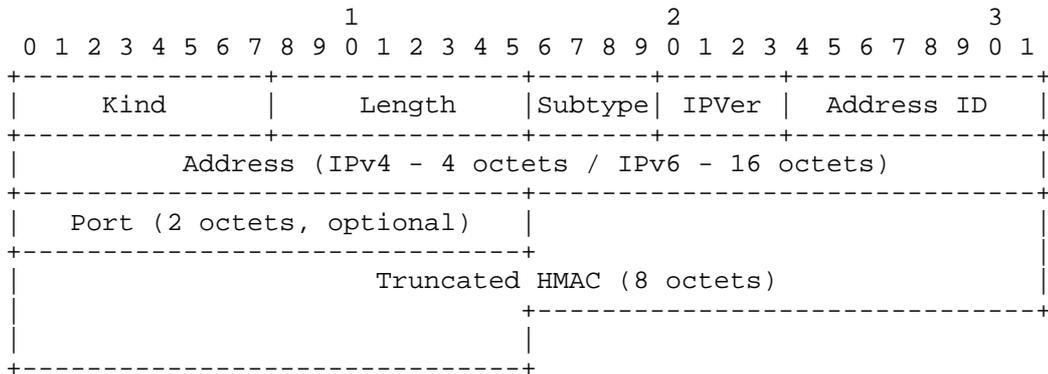


Figure 2: The ADD_ADDR2 option

2.1. Usage of the Address Identifiers

A second issue with the ADD_ADDR option is the management of the address identifiers. At first glance, a Multipath TCP implementation could maintain a table of the IP addresses bound to the local host and associate one identifier with each address. When a new IP address is configured, it is added to the table and the index in the table can be used as its identifier. If a local address stops to be bound to the host, the Multipath TCP can extract its identifier from the table and send the REMOVE_ADDR option over all existing Multipath TCP connections. Unfortunately, such a naive implementation is not possible with the current Multipath TCP implementation.

As defined in [RFC6824], the identifiers 0 are assigned to the addresses that were used for the establishment of the initial subflow. This is because the MP_CAPABLE option does not contain any field to encode an address identifier in contrast with the MP_JOIN option.

An annoying consequence of this design choice is that a Multipath TCP implementation must at least remember the identifier of the address that was used to create the initial subflow. It cannot simply rely on the global address table described above because when an address fails, it must be able to send a REMOVE_ADDR with for address identifier 0 if this address was used to create the initial subflow. This forces a Multipath TCP implementation to at least store the address identifier of the initial subflow for each connection.

One suggestion to ease the maintenance of the addresses on a Multipath TCP implementation would be to stop assuming that the address identifier 0 corresponds to the address used to establish the initial subflow. Instead, the implementation should maintain a table of all the addresses that it uses with Multipath TCP and assign one strictly positive identifier to each address. In this case, each address assigned to the host has the same address identifier for all the Multipath TCP connections. When a new address is learned, it is automatically assigned the next available address identifier and can be announced over all existing Multipath TCP connections depending on the policy applied for the address announcements. When an address is not bound anymore to this host, then the same REMOVE_ADDR option can be sent over all Multipath TCP connections.

There is one missing element in the solution discussed above : how to announce the real address identifier that corresponds to the initial subflow. A simple solution to this problem is to use the ADD_ADDR option without an address as shown in Figure 3.

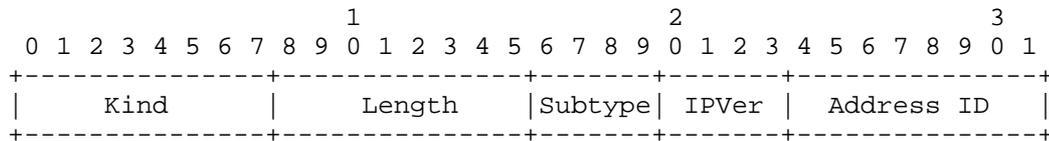


Figure 3: The ADD_ADDR option indicating the address identifier of the current subflow

This empty ADD_ADDR option indicates the address identifier of the address on the subflow over which it has been transmitted. It MUST only be used on the initial subflow since on the other subflows the same information is carried reliably in the MP_JOIN option of the SYN segments. The IPVer field of this ADD_ADDR option MUST match the IP version used for the initial subflow.

2.2. Reliability of the ADD_ADDR Option

A third issue with the ADD_ADDR option is that since it is transmitted as a TCP option, it is not delivered reliably [Cellnet12]. If it announces an IPv4 address, the ADD_ADDR option could be inserted inside a segment that carries data and would thus be delivered reliably like the user data. However, if the ADD_ADDR option contains an IPv6 address, it might be too large to fit inside a segment that already contains a DSS option and possibly other options such as the [RFC1323] timestamps. Given its length, the ADD_ADDR2 option cannot be placed in the same segment as a DSS option. In these two cases, the ADD_ADDR/ADD_ADDR2 option will be often transmitted inside a duplicate ACK that is not delivered reliably. [Cellnet12] proposes a method to improve the reliability of the transmission of the ADD_ADDR option, but to our knowledge this method has never been implemented. To cope with packet losses, some implementations could decide to transmit several copies of the ADD_ADDR option over the same connection.

3. Learning the Addresses bound to a host through the DNS

[RFC6824] defines the ADD_ADDR option as the basic technique to learn the addresses bound to the remote host. Given the importance of learning those addresses, one would expect this technique to be supported by all Multipath TCP implementations. This is not the case, since only the Linux implementation of Multipath TCP supports the ADD_ADDR option [I-D.eardley-mptcp-implementations-survey] as defined in [RFC6824]. The other implementations do not support this option [I-D.eardley-mptcp-implementations-survey]. This design choice was probably motivated by security concerns with this option and also because these implementations assume that only the client

creates the subflows and the server is single-homed. In this case, the client (e.g. a smartphone), can create the subflows from any of its own addresses towards the single address of the server.

However, with the deployment of IPv6, the number of dual-stack clients and servers will grow and it will be important for a host that creates a connection towards the IPv4 address of a server to also learn the IPv6 address associated to this particular server. We show in this section that the DNS could be used to distribute the addressing information that is required by Multipath TCP.

There are three possibilities to use to DNS to distribute the list of addresses associated to a given server. A first approach is to use the existing forward DNS and consider that all the 'A' and 'AAAA' records associated with a name correspond to the same server and can be used to establish Multipath TCP subflows. Unfortunately, when several records are associated to a DNS name, this is often for load balancing reasons and those records point to the addresses of different hosts. A second approach would be to define a new DNS record that contains the list of the IP addresses associated to a given host. However, this would require to deploy a new type of DNS record. Proposals that were made in the past to define new RR types were not endorsed by the IETF (e.g., one single RR for dual stack hosts [I-D.li-dnsex-ipv4-ipv6] or a distinct RR for IPv4-Embedded IPv6 Address [I-D.boucadair-behave-dns-a64]).

The third approach that we propose in this document is to use the reverse DNS to encode the information about the alternate addresses that are associated to a given host. The reverse DNS tree typically only contains PTR records that associate names to reverse IPv4 or IPv6 addresses. However, nothing prevents the use of the reverse DNS to store A and AAAA records. This is the approach that we recommend. It does not require any change to the DNS protocol and can leverage dynamic updates to the DNS [RFC3007] and DNSSEC to authenticate the advertisement of addresses [RFC4034].

As an example, consider the server whose name is `mptcp.example.org` and which is reachable via the following IP addresses taken from the documentation prefixes [RFC3849] [RFC5737] :

- o 192.0.2.10
- o 198.51.100.23
- o 2001:db8::1234

The forward DNS will contain the following records for this server

- [I-D.boucadair-behave-dns-a64]
Boucadair, M. and E. Burgey, "A64: DNS Resource Record for IPv4-Embedded IPv6 Address", draft-boucadair-behave-dns-a64-02 (work in progress), September 2010.
- [I-D.eardley-mptcp-implementations-survey]
Eardley, P., "Survey of MPTCP Implementations", draft-eardley-mptcp-implementations-survey-02 (work in progress), July 2013.
- [I-D.ietf-mptcp-rfc6824bis]
Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", draft-ietf-mptcp-rfc6824bis-04 (work in progress), March 2015.
- [I-D.li-dnsexp-ipv4-ipv6]
Li, L., Li, Z., and X. Duan, "DNS Extensions to Support IPv4 and IPv6", draft-li-dnsexp-ipv4-ipv6-02 (work in progress), October 2009.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC1323] Jacobson, V., Braden, B., and D. Borman, "TCP Extensions for High Performance", RFC 1323, May 1992.
- [RFC3007] Wellington, B., "Secure Domain Name System (DNS) Dynamic Update", RFC 3007, November 2000.
- [RFC3849] Huston, G., Lord, A., and P. Smith, "IPv6 Address Prefix Reserved for Documentation", RFC 3849, July 2004.
- [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.
- [RFC5737] Arkko, J., Cotton, M., and L. Vegoda, "IPv4 Address Blocks Reserved for Documentation", RFC 5737, January 2010.
- [RFC6182] Ford, A., Raiciu, C., Handley, M., Barre, S., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development", RFC 6182, March 2011.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013.

[TMA2015] Hesmans, B., Hoang Tran-Viet, ., Sadre, R., and O. Bonaventure, "A first look at real Multipath TCP traffic", TMA 2015 , April 2015, <<http://inl.info.ucl.ac.be/publications/first-look-real-multipath-tcp-traffic>>.

Author's Address

Olivier Bonaventure
UCLouvain

Email: Olivier.Bonaventure@uclouvain.be

MPTCP Working Group
Internet-Draft
Updates: 6824 (if approved)
Intended status: Experimental
Expires: December 18, 2015

O. Bonaventure
B. Hesmans
UCLouvain
M. Boucadair
Orange
June 16, 2015

Experimental Multipath TCP option
draft-bonaventure-mptcp-exp-option-00

Abstract

This document proposes a generic format for an experimental Multipath TCP option. This experimental option is expected to replace the option type that was reserved for private use in [RFC6824].

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 December 18, 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

1. Introduction	2
2. The Experimental Multipath TCP option	3
3. Security considerations	6
4. IANA considerations	6
5. Conclusion	6
6. Acknowledgements	6
7. References	7
7.1. Normative References	7
7.2. Informative References	7
Authors' Addresses	8

1. Introduction

Multipath TCP is an extension to TCP [RFC0793] that was specified in [RFC6824]. Like other TCP extensions, Multipath TCP relies on TCP options to exchange information. To mitigate some forms of middlebox interference, Multipath TCP uses only one TCP option identified by Kind=30 [RFC6824]. TCP Option Kind=30 has a variable length and includes a subtype field that can support up to 16 different Multipath TCP options. As of this writing, 8 different Multipath TCP options are defined [RFC6824] :

- o subtype=0x0 : MP_CAPABLE
- o subtype=0x1 : MP_JOIN
- o subtype=0x2 : DSS
- o subtype=0x3 : ADD_ADDR
- o subtype=0x4 : REMOVE_ADDR
- o subtype=0x5 : MP_PRIO
- o subtype=0x6 : MP_FAIL
- o subtype=0x7 : MP_FASTCLOSE

In addition, subtype=0x8 is already reserved for a new version of the ADD_ADDR option [I-D.ietf-mptcp-rfc6824bis] and subtype=0xf is reserved for use in private testbeds [RFC6824]. There are thus less than six different Multipath TCP options that are available to extend the protocol.

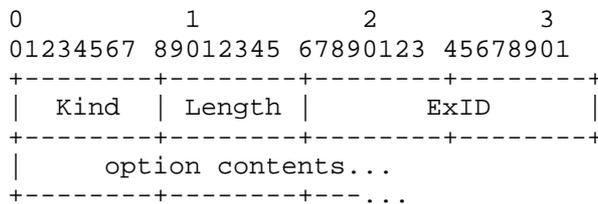
The Multipath TCP options defined above cover the basic functionalities for Multipath TCP but previous experience with TCP

indicates that it can be expected that Multipath TCP implementors will propose new extensions that require a new option subtype. Some authors have already proposed such extensions [I-D.boucadair-mptcp-symmetric] [I-D.paasch-mptcp-syncookies] [I-D.xu-mptcp-prmp]. Given the small number of available Multipath TCP options, it will be difficult to officially reserve them for all proposed utilisations even if experiments with these new Multipath TCP options would probably be useful.

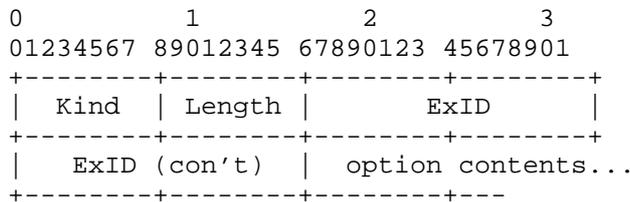
This document proposes a shared use of the Multipath TCP option subtype=0xf following a similar approach as described in [RFC6994] for TCP. This approach is efficient as it encourages conducting new experiments while not consuming the limited option codepoint space.

2. The Experimental Multipath TCP option

The extensibility of TCP depends largely on the definition of new TCP options. As of this writing, 30 different TCP options have been assigned by the IANA among the 256 possible values. Two of these option types (253 and 254) have been reserved as experimental options in [RFC4727]. However, this document did not specify how these experimental options should be used. [RFC6994] clarified the situation by proposing two encodings for these experimental options to support a shared use of the same option kind by different experimental options. Basically, [RFC6994] proposes to use a 16 bits experiment identifier in the first two bytes of the option as shown in the figure below.

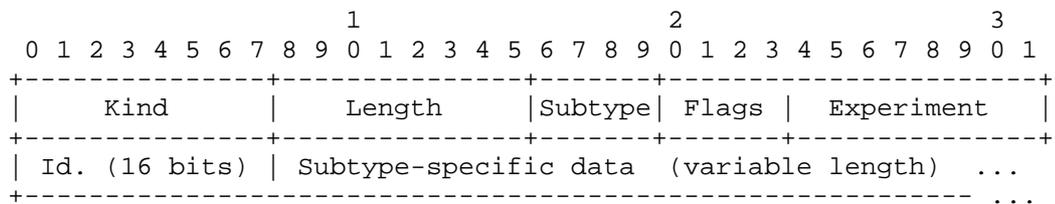


[RFC6994] extends this format to support 32 bits identifiers where the second 16 bits word is used to encode a magic number.

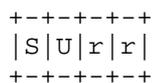


The experimental TCP options defined in [RFC6994] could in theory be used over any TCP connection, including the subflows used by Multipath TCP. However, using them as such to support experimental features on a Multipath TCP connection would cause two practical problems. The first problem is that the utilisation of such an option should be negotiated during the three-way handshake by adding the experimental option in the SYN and SYN+ACK segments. Given the limited space for the options in these segments when Multipath TCP is used and other options are already used, this is not a valid approach. Second, middleboxes could accept the Multipath TCP option and drop or modify an experimental TCP option that is used to support an experimental Multipath TCP feature. This would increase the complexity of the middlebox handling algorithms inside a Multipath TCP implementation.

For the above two reasons, we propose the utilisation of an Experimental Multipath TCP option that is encoded as shown in the figure below.



Flags is a set of 4 flags:



where "rr" are for future assignment as additional flag bits. r bits MUST each be sent as zero and MUST be ignored on receipt.

The Experimental Multipath TCP option has a variable length. It is always of kind=30 and Subtype=0xf. The Experiment identifier is a 16 bits integer that shall be assigned by using the same procedure as defined in [RFC6994]. [RFC6994] also supports the utilisation of 32 bits Experiment Identifiers. These 32 bits identifiers are composed of a 16 bits identifier assigned on a FCFS basis and a 16 bits magic number. At this stage, we propose that the Experimental Multipath TCP option uses only the 16 bits Experiment Identifier.

The two high order flags that are included in the Experimental Multipath TCP option have the following semantics :

- o "S" flag (highest order bit) : This is the synchronising bit. When set to 1, it indicates that the host sending this option expects a reply from the remote host with an option having the same experiment identifier, but possibly containing other data.
- o "U" flag (second highest order bit) : When set to 1, this flag indicates that the experimental option was received by the sending host but it was unable to parse it.

The two low order flags are currently reserved for further use. They MUST be set to zero when sending and ignored upon reception.

To use the Experimental Multipath TCP option with a given experiment identifier over a Multipath TCP connection, the sending host must first verify the ability of the remote host to support this particular Experimental option. For this, it first sends in any valid TCP segment, including a duplicate acknowledgement, an Experimental Multipath TCP option with the "S" flag set. Upon reception of this option, the receiving host will verify whether it supports it. If yes, it shall return a TCP segment that contains the experimental option with the same identifier and the "S" and the "U" flags reset. This option may contain additional data depending on its semantics. If the receiving host does not recognise the experimental option that it has received, it shall return a TCP segment that contains the received experimental option with the "S" flag set to 0 and the "U" flag set to 1. Upon reception of this option, the sending host knows that the experimental option cannot be used over this Multipath TCP connection.

Once the utilisation of a specific Experimental Multipath TCP option has been negotiated over a Multipath TCP connection, both hosts can use this experimental Multipath TCP option. If a host receives an Experimental Multipath TCP option with the "U" flag set to 0 that it does not support or contains information that it cannot parse, it shall return the exact option that it received with the "U" flag set to indicate the error to the remote host. If an invalid option is received with the "U" flag set to 0, it must be silently discarded.

Future documents specifying new experimental MPTCP options should specify the exact semantic of the Subtype-specific data and whether additional validation operations are to be followed at both sides. Note, data can be included in an experimental data concurrently with the capability check (S/U).

As a side effect of the use of a shared codepoint, an overhead of the 2 octets will be consumed in the option space for each experimental MPTCP option. This overhead should be taken into account when drawing conclusions related to the co-existence of multiple MPTCP

experimental options or, even worse, between experimental MPTCP options and experimental TCP options.

3. Security considerations

This document does not modify the security considerations for Multipath TCP.

4. IANA considerations

This document proposes to modify the Multipath TCP option subtype 0xf that was reserved for utilisation in private testbed to support the Experimental Multipath TCP option.

The same considerations as for [RFC6994] apply. IANA should create a "Multipath TCP Experimental Option Identifiers (MPTCP ExIDs)" registry. This registry contains the 16 bits ExIDs and a reference (description, document pointer, or assignee name and e-mail contact) for each entry. MPTCP ExIDs are assigned on a First Come, First Served (FCFS) [RFC5226].

IANA will advise applicants of duplicate entries to select an alternate value, as per typical FCFS processing.

IANA will record known duplicate uses to assist the community in both debugging assigned uses as well as correcting unauthorized duplicate uses.

IANA should impose no requirement on making a registration other than indicating the desired codepoint and providing a point of contact. A short description or acronym for the use is desired but should not be required.

5. Conclusion

This document has proposed a new experimental Multipath TCP option that is intended to replace the already defined, but apparently currently unused, Multipath TCP option reserved for private testbed. This new Experimental Multipath TCP option will enable implementors to perform experiments with Multipath TCP options without using the limited Multipath TCP option subtypes.

6. Acknowledgements

This work was partially supported by the FP7-Trilogy2 project. This document was heavily inspired by [RFC6994].

7. References

7.1. Normative References

- [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", draft-boucadair-mptcp-symmetric-02 (work in progress), March 2015.
- [I-D.ietf-mptcp-rfc6824bis]
Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", draft-ietf-mptcp-rfc6824bis-04 (work in progress), March 2015.
- [I-D.paasch-mptcp-syncookies]
Paasch, C., Biswas, A., and D. Haas, "Making Multipath TCP robust for stateless webservers", draft-paasch-mptcp-syncookies-00 (work in progress), April 2015.
- [I-D.xu-mptcp-prmp]
Xu, C., Huang, H., Zhang, H., Xiong, C., and L. Zhu, "Multipath Transmission Control Protocol (MPTCP) Partial Reliability Extension", draft-xu-mptcp-prmp-00 (work in progress), April 2015.
- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC4727] Fenner, B., "Experimental Values In IPv4, IPv6, ICMPv4, ICMPv6, UDP, and TCP Headers", RFC 4727, November 2006.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC6994] Touch, J., "Shared Use of Experimental TCP Options", RFC 6994, August 2013.

Authors' Addresses

Olivier Bonaventure
UCLouvain

Email: Olivier.Bonaventure@uclouvain.be

Benjamin Hesmans
UCLouvain

Email: Benjamin.Hesmans@uclouvain.be

Mohamed Boucadair
Orange

Email: mohamed.boucadair@orange.com

Network Working Group
Internet-Draft
Updates: 6824 (if approved)
Intended status: Experimental
Expires: September 10, 2015

M. Boucadair
C. Jacquenet
France Telecom
March 9, 2015

An Extension to MPTCP for Symmetrical Sub-Flow Management
draft-boucadair-mptcp-symmetric-02

Abstract

This document specifies a MPTCP extension that allows to achieve symmetrical subflow management. In particular, this extension allows both endpoints to add new subflows whenever needed without waiting for the endpoint which initiated the first subflow to add new ones.

This document updates RFC 6824.

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 10, 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

1. Introduction	2
2. Problem Space	2
3. Proposed Solution	3
4. Security Considerations	5
5. IANA Considerations	5
6. Acknowledgements	5
7. References	6
7.1. Normative References	6
7.2. Informative References	6
Authors' Addresses	7

1. Introduction

This document specifies a MPTCP [RFC6824] extension to achieve symmetrical subflow management. The problem space is further described in Section 2, while a proposed solution is discussed in Section 3.

This document assumes Port Control Protocol (PCP)-enabled networks [RFC6887]. But other procedures can be used to instantiate mappings and discover the external LP address/port assigned by an upstream flow-aware device (e.g., CGN [RFC6888], firewall, etc.).

2. Problem Space

The following is extracted from[I-D.ietf-mptcp-experience]:

From a subflow viewpoint, the Multipath TCP protocol is completely symmetrical. Both the clients and the server have the capability to create subflows. However in practice the existing Multipath TCP implementations [I-D.eardley-mptcp-implementations-survey] have opted for a strategy where only the client creates new subflows. The main motivation for this strategy is that often the client resides behind a NAT or a firewall, preventing passive subflow openings on the client.

This means that in practice only the client (that is the TCP endpoint that initiated the first subflow) can initiate new subflows. This is not optimal in situations where (1) the remote endpoints want to boost their sending rate or handover to a new IP address without waiting for the client to add new subflows, (2) or when the traffic distribution as observed by the remote endpoint does not meet its local policies. Adding new subflows should be subject to both the client's and server's local policies, not only those of the client.

3. Proposed Solution

This procedure can be activated upon bootstrap or when a network attachment change occurs (e.g., attach to a new network); it is not executed for every new MPTCP connection:

- o Each endpoint proceeds to the discovery of upstream flow-aware devices (e.g., NAT, Firewall).

This can be achieved by various means, e.g., using UPnP IGD [IGD1][IGD2], PCP server discovery [RFC6887], PCP DHCP option [RFC7291], DS-Lite AFTR [RFC6334], etc.

A NAT/firewall can be embedded in a CPE (Customer Premises Equipment) and/or hosted in the network provider's side.

- o Appropriate mappings are instantiated in those discovered flow-aware devices. In particular, external IP address(es) and port numbers are retrieved.

This can be achieved using PCP [RFC6887]. Note, mappings created by PCP MAP requests are, by definition, endpoint-independent mappings (EIMs) with endpoint-independent filtering (EIF). Filters can be associated with the PCP MAP request to restrict a mapping to be bound to specific remote peer(s).

PCP allows to dynamically control both NATs and firewalls. Furthermore, PCP allows to retrieve the lifetime associated with an external IP address and external port number.

If the host is a UPnP IGD Control Point, [RFC6970] allows to relay UPnP IGD primitives into PCP messages. PCP can also be used to control multi-layered NATs/firewall owing to the activation of [I-D.ietf-pcp-proxy] in intermediate NATs/firewalls.

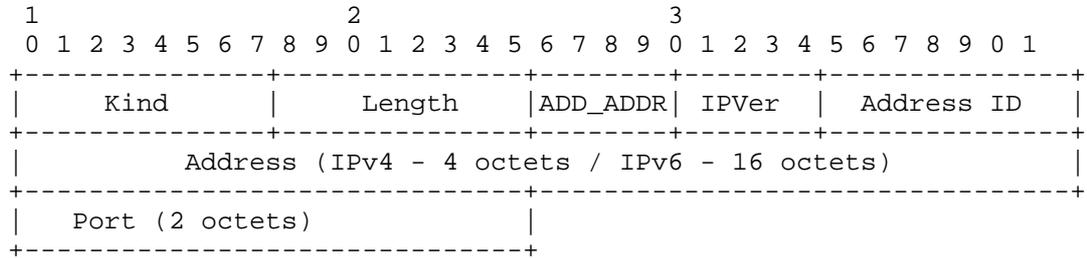
- o When initiating an MPTCP connection, external IP addresses and port numbers are communicated to the remote peer.

This can be achieved using ADD_ADDR together with a new option that will indicate that the address/port pair (identified by Address ID) enclosed in ADD_ADDR has been checked so that incoming flows can be sent to this address/port.

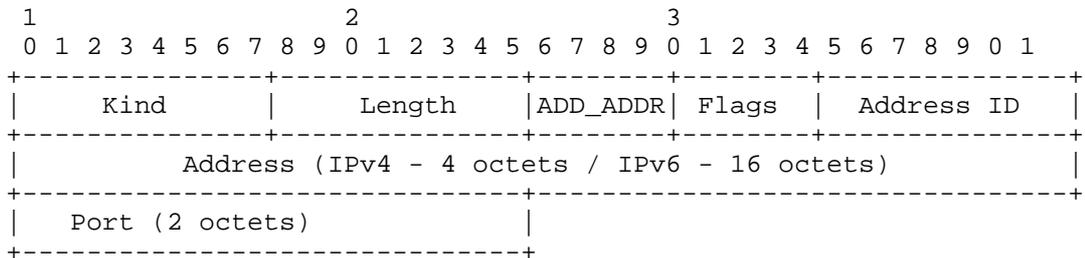
A second implementation flavor is to define a new option, similar to ADD_ADDR, but which will include an optional field (lifetime). The lifetime can be used as an input to the traffic management block at the remote side. This field can be derived from the lease returned in DHCP, or in PCP requests. The use of this option is an indication that appropriate actions were undertaken at the remote side to receive incoming packets. A remote peer can use the enclosed address/port to add a new subflow without any risk to experience failures at the client side. Indicating the lifetime associated with an IP address is seen as a limitation of current APIs as discussed in Section-3.2.1 of [RFC6250]. The lifetime can be used as a hint to migrate flows to another subflows.

Another implementation flavor is to update ADD_ADDR as shown in Figure 1. "IPVer" is useless since the length is sufficient to determine whether the enclosed IP address is IPv4 or IPv6.

OLD:



NEW:



Flags is a set of 4 flags: +---+---+---+
 |C|r|r|r|
 +---+---+---+

C flag MUST be set to 1 when the address/port are checked.
 "rrr" are for future assignment as additional flag bits.
 r bits MUST each be sent as zero and MUST be ignored on receipt.

Figure 1

4. Security Considerations

PCP-related security considerations are discussed in [RFC6887].
 MPTCP-related security considerations are documented in [RFC6824] and [I-D.ietf-mptcp-attacks].

5. IANA Considerations

TBC.

6. Acknowledgements

Many thank to Olivier Bonaventure who suggested the idea of updating ADD_ADDR.

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.eardley-mptcp-implementations-survey] Eardley, P., "Survey of MPTCP Implementations", draft-eardley-mptcp-implementations-survey-02 (work in progress), July 2013.
- [I-D.ietf-mptcp-attacks] 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.
- [I-D.ietf-mptcp-experience] Bonaventure, O., Paasch, C., and G. Detal, "Experience with Multipath TCP", draft-ietf-mptcp-experience-01 (work in progress), March 2015.
- [I-D.ietf-pcp-proxy] Perreault, S., Boucadair, M., Penno, R., Wing, D., and S. Cheshire, "Port Control Protocol (PCP) Proxy Function", draft-ietf-pcp-proxy-06 (work in progress), December 2014.
- [IGD1] UPnP Forum, , "WANIPConnection:1 Service (<http://www.upnp.org/specs/gw/UPnP-gw-WANIPConnection-v1-Service.pdf>)", November 2001.
- [IGD2] UPnP Forum, , "WANIPConnection:2 Service (<http://upnp.org/specs/gw/UPnP-gw-WANIPConnection-v2-Service.pdf>)", September 2010.
- [RFC6250] Thaler, D., "Evolution of the IP Model", RFC 6250, May 2011.
- [RFC6334] Hankins, D. and T. Mrugalski, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite", RFC 6334, August 2011.

- [RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013.
- [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.
- [RFC6970] Boucadair, M., Penno, R., and D. Wing, "Universal Plug and Play (UPnP) Internet Gateway Device - Port Control Protocol Interworking Function (IGD-PCP IWF)", RFC 6970, July 2013.
- [RFC7291] Boucadair, M., Penno, R., and D. Wing, "DHCP Options for the Port Control Protocol (PCP)", RFC 7291, July 2014.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Christian Jacquenet
France Telecom
Rennes 35000
France

Email: christian.jacquenet@orange.com

Network Working Group
Internet-Draft
Intended status: Informational
Expires: June 5, 2020

T. Dreibholz
SimulaMet
S. Ferlin
Simula Research Laboratory
O. Alay
A. Elmokashfi
I. Livadariu
SimulaMet
X. Zhou
Hainan University
December 3, 2019

MPTCP Experiences in the NorNet Testbed
draft-dreibholz-mptcp-nornet-experience-05

Abstract

This document collects some experiences of Multi-Path TCP (MPTCP) evaluations in the NorNet testbed.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on June 5, 2020.

Copyright Notice

Copyright (c) 2019 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 (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
1.1. Abbreviations	2
1.2. Multi-Path TCP	2
1.3. Scope	2
2. Testbed Evaluation	3
2.1. The NorNet Testbed	3
2.2. Multi-Path TCP in NorNet	3
2.3. NetPerfMeter	3
3. Research Results and Work in Progress	4
4. Security Considerations	5
5. IANA Considerations	5
6. Acknowledgments	5
7. References	6
7.1. Normative References	6
7.2. Informative References	6
Authors' Addresses	11

1. Introduction

1.1. Abbreviations

- o TCP: Transmission Control Protocol
- o MPTCP: Multi-Path TCP

1.2. Multi-Path TCP

The Multi-Path TCP (MPTCP) extension for the Transmission Control Protocol (TCP) has been defined in [3], [4], [1], [2], [5]. There are also detailed introductions provided for example by [8], [9] as well as lots of further information material on [6]. MPTCP is therefore not introduced in more detail here.

1.3. Scope

The scope of this document is to collect some experiences with the usage of MPTCP in the NorNet testbed, a large-scale Internet testbed for multi-homed systems.

2. Testbed Evaluation

2.1. The NorNet Testbed

The NorNet testbed (<https://www.nntb.no>) introduced in [11], [15], [10] is a programmable testbed platform with focus on the evaluation of multi-homed systems. It consists of programmable nodes that are distributed all over Norway as well as further locations in other countries. NorNet is built and operated by the Simula Research Laboratory and financed by Forskningsraedet (the Research Council of Norway) through their INFRASTRUKTUR program (project number 208798/F50).

NorNet has two main components: NorNet Core and NorNet Edge. NorNet Core consists of currently 20 programmable sites (<https://www.nntb.no/pub/nornet-configuration/NorNetCore-Sites.html>), most of them multi-homed to several network providers. Details can be found in [12], [14], [11] and [16]. NorNet Edge consists of several hundreds of smaller nodes connected to all mobile broadband providers in Norway (<http://robustenett.no/map>), details can be found in [13]. Together, these two components offer a unique platform for experimental networking research. NorNet is made available to the international networking research community.

Further details on NorNet can be found on the NorNet website [17].

2.2. Multi-Path TCP in NorNet

The NorNet nodes are Linux-based and therefore run the Linux MPTCP implementation by Universite catholique de Louvain-la-Neuve (see [18]). With support by all NorNet Core nodes (circa 100 nodes at currently 20 sites) and most NorNet Edge nodes (several hundreds), NorNet probably provides the world's largest MPTCP experimentation platform.

2.3. NetPerfMeter

NetPerfMeter [19], [7] is a network performance meter for the UDP, TCP, MPTCP, SCTP and DCCP transport protocols over IPv4 and IPv6. It simultaneously transmits bidirectional flows to an endpoint and measures the resulting flow bandwidths and QoS. The results are written as vector and scalar files. NetPerfMeter is provided in NorNet to allow for transport protocol comparisons and transport protocol performance evaluations.

3. Research Results and Work in Progress

So far, work has been done on evaluating MPTCP in the real-world Internet. The following list is a short overview of current research:

- o [20] examines the different properties of 3G paths (UMTS as well as CDMA2000) in NorNet Edge. These results provide an overview of the QoS characteristics that can be expected by multi-path transport protocols (and particularly by MPTCP) on different paths.
- o [21] analyzes the impact of buffer bloat on MPTCP connections over mobile broadband (2G, 3G) and WLAN paths. Furthermore, it proposes and evaluates Multi-Path Transport Bufferbloat Mitigation (MPT-BM), a bufferbloat mitigation algorithm to improve MPTCP performance in buffer-bloated wireless networks.
- o [22] examines the performance benefits of multi-path transport with MPTCP under heterogeneous wireless networks. It furthermore introduces and evaluates the Dynamic Relative Path Scoring (DRePaS) algorithm that optimizes the path management in such setups.
- o [23] examines the performance benefits of the state-of-the-art Linux MPTCP implementation in a large-scale NorNet setup, covering sites in multiple countries on different continents. It particularly also shortly introduces the NetPerfMeter measurement tool that is used to perform the measurements, and particularly its extension for MPTCP.
- o [24] examines how the IPv4/IPv6 identity duality can be utilized with MPTCP in order to improve performance even in case of only a single ISP connection.
- o [25] provides an introduction to the NetPerfMeter tool for MPTCP experiments.
- o [26] presents MPTCP's architecture and multi-path congestion control algorithm concepts. Then, it examines three test scenarios in the NorNet testbed, particularly highlighting the performance difference between using uncoupled and coupled congestion controls in multi-homed, real-world Internet setups.
- o [27] introduces a scanning infrastructure to search for MPTCP-capable hosts in the Internet. In a study, it used the top-1M Alexa servers to test the platform and gain insight about server support for MPTCP.

- o [28] proposes a practical shared bottleneck detection (SBD) algorithm for MPTCP, namely MPTCP-SBD. Through extensive emulations, it is shown that MPTCP-SBD outperforms all currently deployed MPTCP coupled congestion controls by accurately detecting bottlenecks resulting in throughput gains in the absence of shared bottlenecks, while remaining fair to TCP in shared bottlenecks scenarios.
- o [29] proposes a send-window BLocking ESTimation scheduler, BLEST, which aims to minimise head-of-line-blocking in heterogeneous networks. BLEST increases the potential for capacity aggregation by reducing the number of spurious retransmissions.
- o [30] examines path management strategies for MPTCP in real-world, multi-homed Internet setups. Particular goal is to highlight the performance impact of different path management and congestion control settings in such realistic scenarios.
- o [31] analyses the actual buffer size requirements for MPTCP in heterogeneous, multi-homed Internet setups. Particular goal is to provide some guidelines on buffer size requirements for achieving a reasonable throughput, while on the other hand not wasting resources.
- o [32] examines the relationship of buffer size with throughput and congestion control algorithms, based on the statistical predictive modelling method.
- o [33] examines the application of MPTCP for real-time video streaming for future Mobile Edge Computing (MEC) scenarios.

4. Security Considerations

Security considerations on MPTCP are described in [3].

5. IANA Considerations

This document has no actions for IANA.

6. Acknowledgments

The authors would like to thank Hakim Adhari, Fa Fu, Feng Zhou and Kun Wang for discussions and support.

7. References

7.1. Normative References

- [1] Raiciu, C., Handley, M., and D. Wischik, "Coupled Congestion Control for Multipath Transport Protocols", RFC 6356, DOI 10.17487/RFC6356, October 2011, <<https://www.rfc-editor.org/info/rfc6356>>.
- [2] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, DOI 10.17487/RFC6824, January 2013, <<https://www.rfc-editor.org/info/rfc6824>>.

7.2. Informative References

- [3] Bagnulo, M., "Threat Analysis for TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6181, DOI 10.17487/RFC6181, March 2011, <<https://www.rfc-editor.org/info/rfc6181>>.
- [4] Ford, A., Raiciu, C., Handley, M., Barre, S., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development", RFC 6182, DOI 10.17487/RFC6182, March 2011, <<https://www.rfc-editor.org/info/rfc6182>>.
- [5] Scharf, M. and A. Ford, "Multipath TCP (MPTCP) Application Interface Considerations", RFC 6897, DOI 10.17487/RFC6897, March 2013, <<https://www.rfc-editor.org/info/rfc6897>>.
- [6] Dreibholz, T., "Thomas Dreibholz's MPTCP Page", Online: <https://www.uni-due.de/~be0001/mptcp/>, 2016, <<https://www.uni-due.de/~be0001/mptcp/>>.
- [7] Dreibholz, T., "NetPerfMeter - A TCP/MPTCP/UDP/SCTP/DCCP Network Performance Meter Tool", Online: <https://www.uni-due.de/~be0001/netperfmeter/>, 2016, <<https://www.uni-due.de/~be0001/netperfmeter/>>.
- [8] Becke, M., "Revisiting the IETF Multipath Extensions on Transport Layer", November 2014, <http://duepublico.uni-duisburg-essen.de/servlets/DerivateServlet/Derivate-37631/Thesis_online.pdf>.

- [9] Ruengeler, I., "SCTP - Evaluating, Improving and Extending the Protocol for Broader Deployment", December 2009, <<http://duepublico.uni-duisburg-essen.de/servlets/DerivateServlet/Derivate-23465/Diss.pdf>>.
- [10] Dreibholz, T., "The NorNet Testbed for Multi-Homed Systems - Introduction and Status", Invited Talk at Princeton University, Department of Computer Science, May 2014, <<https://www.simula.no/sites/www.simula.no/files/publications/Simula.simula.2730.pdf>>.
- [11] Dreibholz, T., "NorNet at NICTA - An Introduction to the NorNet Testbed", Invited Talk at National Information Communications Technology Australia (NICTA), January 2016, <<https://www.simula.no/file/nicta2016-presentationpdf/download>>.
- [12] Gran, E., Dreibholz, T., and A. Kvalbein, "NorNet Core - A Multi-Homed Research Testbed", Computer Networks, Special Issue on Future Internet Testbeds Volume 61, Pages 75-87, ISSN 1389-1286, DOI 10.1016/j.bjp.2013.12.035, March 2014, <<https://www.simula.no/file/simulasimula2236pdf/download>>.
- [13] Kvalbein, A., Baltrunas, D., Evensen, K., Xiang, J., Elmokashfi, A., and S. Ferlin, "The NorNet Edge Platform for Mobile Broadband Measurements", Computer Networks, Special Issue on Future Internet Testbeds Volume 61, Pages 88-101, ISSN 1389-1286, DOI 10.1016/j.bjp.2013.12.036, March 2014, <<https://www.simula.no/sites/www.simula.no/files/publications/Simula.simula.2434.pdf>>.
- [14] Dreibholz, T. and E. Gran, "Design and Implementation of the NorNet Core Research Testbed for Multi-Homed Systems", Proceedings of the 3rd International Workshop on Protocols and Applications with Multi-Homing Support (PAMS) Pages 1094-1100, ISBN 978-0-7695-4952-1, DOI 10.1109/WAINA.2013.71, March 2013, <<https://www.simula.no/sites/www.simula.no/files/publications/threfereedinproceedingsreference.2012-12-20.7643198512.pdf>>.

- [15] Dreibholz, T., "Multi-Path Transport - From Simulations to Real-World Internet Measurements", Keynote Talk at the Universidad de Castilla-La Mancha, Instituto de Investigacion Informatica de Albacete, February 2017, <<https://www.simula.no/file/uclm2017-multipath-presentationpdf-1/download>>.
- [16] Dreibholz, T., "An Experiment Tutorial for the NorNet Core Testbed at the the Universidad de Castilla-La Mancha", Tutorial at the Universidad de Castilla-La Mancha, Instituto de Investigacion Informatica de Albacete, February 2017, <<https://www.simula.no/file/uclm2017-nornet-tutorialpdf/download>>.
- [17] Dreibholz, T., "NorNet - A Real-World, Large-Scale Multi-Homing Testbed", Online: <https://www.nntb.no/>, 2017, <<https://www.nntb.no/>>.
- [18] Paasch, C. and S. Barre, "Multipath TCP implementation in the Linux kernel", Online: <http://www.multipath-tcp.org>, 2016, <<http://www.multipath-tcp.org>>.
- [19] Dreibholz, T., Becke, M., Adhari, H., and E. Rathgeb, "Evaluation of A New Multipath Congestion Control Scheme using the NetPerfMeter Tool-Chain", Proceedings of the 19th IEEE International Conference on Software, Telecommunications and Computer Networks (SoftCOM) Pages 1-6, ISBN 978-953-290-027-9, September 2011, <<https://www.wiwi.uni-due.de/fileadmin/fileupload/I-TDR/SCTP/Paper/SoftCOM2011.pdf>>.
- [20] Ferlin, S., Dreibholz, T., Alay, Oe., and A. Kvalbein, "Measuring the QoS Characteristics of Operational 3G Mobile Broadband Networks", Proceedings of the 4th International Workshop on Protocols and Applications with Multi-Homing Support (PAMS) Pages 753-758, ISBN 978-1-4799-2652-7, DOI 10.1109/WAINA.2014.123, May 2014, <<https://www.simula.no/sites/www.simula.no/files/publications/Simula.simula.2516.pdf>>.
- [21] Ferlin, S., Dreibholz, T., and Oe. Alay, "Tackling the Challenge of Bufferbloat in Multi-Path Transport over Heterogeneous Wireless Networks", Proceedings of the IEEE/ACM International Symposium on Quality of Service (IWQoS) Pages 123-128, ISBN 978-1-4799-4852-9, DOI 10.1109/IWQoS.2014.6914310, May 2014, <<https://www.simula.no/sites/www.simula.no/files/publications/Simula.simula.2722.pdf>>.

- [22] Ferlin, S., Dreibholz, T., and Oe. Alay, "Multi-Path Transport over Heterogeneous Wireless Networks: Does it really pay off?", Proceedings of the IEEE Global Communications Conference (GLOBECOM) Pages 5005-5011, ISBN 978-1-4799-3512-3, December 2014, <<https://www.simula.no/sites/www.simula.no/files/publications/Simula.simula.2884.pdf>>.
- [23] Dreibholz, T., Zhou, X., and F. Fa, "Multi-Path TCP in Real-World Setups - An Evaluation in the NorNet Core Testbed", 5th International Workshop on Protocols and Applications with Multi-Homing Support (PAMS), March 2015, <<https://www.simula.no/sites/www.simula.no/files/publications/files/pams2015-mptcp-web.pdf>>.
- [24] Livadariu, I., Ferlin, S., Alay, Oe., Dreibholz, T., Dhamdhere, A., and A. Elmokashfi, "Leveraging the IPv4/IPv6 Identity Duality by using Multi-Path Transport", Proceedings of the 18th IEEE Global Internet Symposium (GI), April 2015, <https://www.simula.no/sites/www.simula.no/files/publications/files/gis2015_0.pdf>.
- [25] Dreibholz, T., "NetPerfMeter: A Network Performance Metering Tool", Multipath TCP Blog, September 2015, <<http://blog.multipath-tcp.org/blog/html/2015/09/07/netperfmeter.html>>.
- [26] Fa, F., Zhou, X., Dreibholz, T., Wang, K., Zhou, F., and Q. Gan, "Performance Comparison of Congestion Control Strategies for Multi-Path TCP in the NorNet Testbed", Proceedings of the 4th IEEE/CIC International Conference on Communications in China (ICCC) Pages 607-612, ISBN 978-1-5090-0243-6, November 2015, <<https://www.simula.no/sites/www.simula.no/files/publications/files/fufa-mptcp-web.pdf>>.
- [27] Mehani, O., Holz, R., Ferlin, S., and R. Boreli, "An Early Look at Multipath TCP Deployment in the Wild", Proceedings of the 6th International Workshop on Hot Topics in Planet-Scale Measurement ISBN 978-1-4503-3534-8, DOI 10.1145/2798087.2798088, 2015, <<https://www.nicta.com.au/pub?doc=8791?filename=Mehani2015.pdf>>.

- [28] Ferlin, S., Alay, Oe., Hayes, D., Dreibholz, T., and M. Welzl, "Revisiting Congestion Control for Multipath TCP with Shared Bottleneck Detection", Proceedings of the 35th IEEE International Conference on Computer Communications (INFOCOM), April 2016.
- [29] Ferlin, S., Alay, Oe., Mehani, O., and R. Boreli, "BLEST: Blocking Estimation-Based MPTCP Scheduler for Heterogeneous Networks", Proceedings of the IFIP Networking Conference Pages 431-439, DOI 10.1109/IFIPNetworking.2016.7497206, May 2016, <<http://dl.ifip.org/db/conf/networking/networking2016/1570234725.pdf>>.
- [30] Wang, K., Dreibholz, T., Zhou, X., Fu, F., Tan, Y., Cheng, X., and Q. Tan, "On the Path Management of Multi-Path TCP in Internet Scenarios based on the NorNet Testbed", Proceedings of the IEEE International Conference on Advanced Information Networking and Applications (AINA) ISBN 978-1-5090-6028-3, March 2017, <<https://www.simula.no/file/kunwang-paper-webpdf/download>>.
- [31] Zhou, F., Dreibholz, T., Zhou, X., Fu, F., Tan, Y., and Q. Gan, "The Performance Impact of Buffer Sizes for Multi-Path TCP in Internet Setups", Proceedings of the IEEE International Conference on Advanced Information Networking and Applications (AINA) ISBN 978-1-5090-6028-3, March 2017, <<https://www.simula.no/file/fengzhou-mptcp-webpdf/download>>.
- [32] Tan, Q., Yang, X., Zhao, L., Zhou, X., and T. Dreibholz, "A Statistic Procedure to Find Formulae for Buffer Size in MPTCP", Proceedings of the 3rd IEEE Advanced Information Technology, Electronic and Automation Control Conference (IAEAC) Pages 900-907, ISBN 978-1-5386-4509-3, October 2018, <<https://www.simula.no/file/iaeac2018-webpdf/download>>.
- [33] Luo, Y., Zhou, X., Dreibholz, T., and H. Kuang, "A Real-Time Video Streaming System over IPv6+MPTCP Technology", Proceedings of the 1st International Workshop on Recent Advances for Multi-Clouds and Mobile Edge Computing (M2EC) in conjunction with the 33rd International Conference on Advanced Information Networking and Applications (AINA), March 2019, <<https://www.simula.no/file/m2ec2019-videostreamingpdf/download>>.

Authors' Addresses

Thomas Dreibholz
Simula Metropolitan Centre for Digital Engineering
Pilestredet 52
0167 Oslo, Oslo
Norway

Phone: +47-6782-8200
Fax: +47-6782-8201
Email: dreibh@simula.no
URI: <https://www.simula.no/people/dreibh>

Simone Ferlin
Simula Research Laboratory, Networks Department
Martin Linges vei 17
1364 Fornebu, Akershus
Norway

Phone: +47-4072-0702
Fax: +47-4072-0702
Email: ferlin@simula.no
URI: <https://www.simula.no/people/ferlin>

Ozgu Alay
Simula Metropolitan Centre for Digital Engineering
Pilestredet 52
0167 Oslo, Oslo
Norway

Phone: +47-9848-5362
Fax: +47-9848-5362
Email: ozgu@simula.no
URI: <https://www.simula.no/people/ozgu>

Ahmed Elmokashfi
Simula Metropolitan Centre for Digital Engineering
Pilestredet 52
0167 Oslo, Oslo
Norway

Phone: +47-474-52-315
Fax: +47-6782-8201
Email: ahmed@simula.no
URI: <https://www.simula.no/people/ahmed>

Ioana Alexandrina Livadariu
Simula Metropolitan Centre for Digital Engineering
Pilestredet 52
0167 Oslo, Oslo
Norway

Phone: +47-453-98-686
Fax: +47-6782-8201
Email: ioana@simula.no
URI: <https://www.simula.no/people/ioana>

Xing Zhou
Hainan University, College of Information Science and Technology
Renmin Avenue 58
570228 Haikou, Hainan
China

Phone: +86-898-66279141
Email: zhouxing@hainu.edu.cn
URI: <http://www.hainu.edu.cn/stm/xinxi/2011330/10409758.shtml>

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 13, 2014

P. Eardley
BT
July 12, 2013

Survey of MPTCP Implementations
draft-eardley-mptcp-implementations-survey-02

Abstract

This document presents results from the survey to gather information from people who have implemented MPTCP, in particular to help progress the protocol from Experimental to Standards track.

The document currently includes answers from four teams: a Linux implementation from UCLouvain, a FreeBSD implementation from Swinburne, an anonymous implementation in a commercial OS, and a NetScaler Firmware implementation from Citrix Systems, Inc. Thank-you!

In summary, we have four independent implementations of all the MPTCP signalling messages, with the exception of address management, and some interoperability testing has been done by the other three implementations with the 'reference' Linux implementation. So it appears that the RFC is (at least largely) clear and correct. On address management, we have only one implementation of ADD_ADDR with two teams choosing not to implement it. We have one implementation of the working group's coupled congestion control (RFC6356) and none of the MPTCP-aware API (RFC6897).

The main suggested improvements are around

- o how MPTCP falls back (if the signalling is interrupted by a middlebox): (1) corner cases that are not handled properly, (2) at the IETF, the MPTCP community should work with middlebox vendors, either to reduce or eliminate the need for fallback or to understand the middlebox interactions better.
- o security: both better MPTCP security (perhaps building on SSL) and a lighter weight mechanism, preferably both in one mechanism.

It is hoped that the next version can include information from any other implementations. If you are an implementer and want to contribute your answers, please see the -01 version of this document for a blank survey ready to be filled in.

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 January 13, 2014.

Copyright Notice

Copyright (c) 2013 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. Survey - summary of replies	4
3. Interesting aspects of replies	6
3.1. Question 1: Your details	6
3.2. Question 2: Preliminary information about your implementation	7
3.3. Question 3: Support for MPTCP's Signalling Functionality	7
3.4. Question 4: Fallback from MPTCP	7
3.5. Question 5: Heuristics	8
3.6. Question 6: Security	9
3.7. Question 7: IANA	9
3.8. Question 8: Congestion control and subflow policy	9
3.9. Question 9: API	10
3.10. Question 10: Deployments, use cases and operational experiences	10
3.11. Question 11: Improvements to RFC6824	11
4. IANA Considerations	11
5. Security Considerations	11
6. Acknowledgements	11
7. Full survey response for Implementation 1	11
8. Full survey response for Implementation 2	19
9. Full survey response for Implementation 3	23
10. Full survey response for Implementation 4	31
11. Normative References	38
Author's Address	38

1. Introduction

The document reports the results from a survey to gather information from people who have implemented MPTCP. The goal is to help progress the protocol from Experimental to Standards track.

Four responses have been received. Thank-you! They are independent implementations:

- o the Linux implementation from UCLouvain,
- o the FreeBSD implementation from Swinburne
- o an anonymous implementation in a commercial OS
- o a NetScaler Firmware implementation from Citrix Systems, Inc.

The Table below presents a highly-compressed summary, with each row corresponding to one question or sub-question of the survey. The following section highlights some interesting aspects of the replies in less compressed form. The full survey responses are in Appendix A, B, C and D.

It is hoped that the next version of this document can include information about a further (independent) implementation:

- o Georg Hampel's user-space implementation (publicly available but not longer maintained)
- o any other implementations.

2. Survey - summary of replies

The Table below presents a highly-compressed summary, with each row corresponding to one question or sub-question of the survey. A column is left blank for any future responses.

Institution	1 UCLouvain	2 Swinburne	3 Anon	4 Citrix
Question 2 asks about some preliminary topics, including whether the implementation is publicly available and interoperability with the Linux implementation (#1).				
OS	Linux	FreeBSD-10	Commercial	NetScaler
v4 & v6	Both	IPv4	Both	Both

public	Yes	Yes	No	Yes (pay)
independent	Yes	Yes	Yes	Yes
interop	Yes(!)	Mostly	Mostly	Yes

Question 3: Support for MPTCP's signalling functionality
MPTCP's signalling messages are: MP_CAPABLE, MP_JOIN, Data transfer (DSS), ADD_ADDR, REMOVE_ADDR, MP_FASTCLOSE. There are sub-questions for MP_JOIN and DSS.

	UCLouvain	Swinburne	Anon	Citrix
MP_CAPABLE	Yes	Yes	Yes	Yes
MP_JOIN	Yes	Yes	Yes	Yes
initiated by	first end	either end	first end	first end
#subflows	32	8	no limit	6
DSS	Yes	Yes	Yes	Yes
DATA ACK	4 bytes	4 or 8 byte	4 or 8 byte	4 or 8 byte
Data seq num	4 bytes	4 or 8 byte	4 or 8 byte	4 or 8 byte
DATA_FIN	Yes	Yes	Yes	Yes
Checksum	Yes	No	Yes	Yes
ADD_ADDR	Yes	No	No (never)	No (never?)
REMOVE_ADDR	Yes	No	Partly	Yes
FAST_CLOSE	Yes	No	Yes	Yes

Question 4 asks about fallback from MPTCP: if a middlebox mangles MPTCP's signalling by removing MP_CAPABLE, MP_JOIN, DSS or DATA_ACK; if data is protected with Checksum in DSS option; if fallback to TCP uses an infinite mapping; and if any corner cases have been found.

	UCLouvain	Swinburne	Anon	Citrix
MP_CAPABLE	Yes	Yes	Yes	Yes
MP_JOIN	Yes	Yes	Yes	Yes
DSS	Yes	No	Yes	Yes
DATA_ACK	Yes	No	No	
Checksum	Yes	No	Yes	Yes
infinite map	Yes	Yes	Yes	Yes
corner cases	No		Yes	Yes

Question 5 asks about heuristics: aspects that are not required for protocol correctness but impact the performance. Questions are about sized the receiver and sender buffers, re-transmission policy, if additional subflows use the same port number as for the first subflow

	UCLouvain	Swinburne	Anon	Citrix
Recv buffer	auto-tune	TCP_MAXWIN	no tuning	tuned
Sendr buffer	auto-tune	cwnd	no tuning	as TCP
Re-transmits	2nd subflow	2nd subflow	2nd subflow	1st subflow
Port usage	same ports	same ports	diff local	

Question 6 asks about what security mechanisms are implemented: the one defined in RFC6824 and any others.

	UCLouvain	Swinburne	Anon	Citrix

HMAC-SHA1	Yes	Yes	Yes	Yes	
other	Yes	No	No	No	
Question 7 asks whether the implementation follows the IANA-related definitions (for TCP Option Kind and sub-registries).					
	UCLouvain	Swinburne	Anon	Citrix	
RFC6824	Yes	Yes	Yes	Yes	
Question 8 asks about congestion control and related issues: how traffic is shared across multiple subflows; support for 'handover'; and support of RFC6356 (or other) coupled congestion control.					
	UCLouvain	Swinburne	Anon	Citrix	
sharing	shared, RTT	shared	active/back	active/back	
handover	Yes		Yes	Yes	
coupled cc	Yes	No	No	No	
other ccc	Yes, OLIA	No	No	No	
MP-PRIO & B	Yes	No	Yes	Yes	
Question 9 is about the API: how legacy applications interact with the MPTCP stack, and if implemented the RFC6897 API for MPTCP-aware applications.					
	UCLouvain	Swinburne	Anon	Citrix	
legacy apps	default	sysctl	private API	configured	
MPTCP API	No	No	No	No	
advanced API	No	No	No	No	
Question 10 gathers some limited information about operational experiences and deployments.					
	UCLouvain	Swinburne	Anon	Citrix	
Scenario	several	several	mobile	proxy	
environment	internet	controlled	internet	internet	
ends / proxy	end hosts	end hosts	end hosts	proxy	

3. Interesting aspects of replies

This section tries to highlight some interesting comments made in the surveys. The Appendices can be consulted for further details.

3.1. Question 1: Your details

Implementation 1 has been implemented by Sebastien Barre, Christoph Paasch and a large team, mainly at UCLouvain. Implementation 2 has been implemented by Lawrence Stewart and Nigel Williams at Swinburne University of Technology. Both these implementations are publicly available. Implementation 3 comes from an anonymous team with a

commercial OS. Implementation 4 comes from Citrix Systems, Inc.

3.2. Question 2: Preliminary information about your implementation

Three of the four implementations are publicly available, two for free (under GPLv2 and BSD licences) and one for a fee (NetScaler Firmware). Implementation 3 (commercial OS) is planned for use in a mobile environment, with MPTCP is used in active/backup mode.

All implementations support IPv4 and three of four support IPv6.

All implementations are being actively worked on, in order to improve performance and stability and conformance with the RFC.

3.3. Question 3: Support for MPTCP's Signalling Functionality

Three of the four implementations have implemented all the MPTCP signalling, with the interesting exception of address management, whilst Implementation 2 plans to add support for all those signalling capabilities it does not yet support.

On address management, two implementations have decided not to implement ADD_ADDR. (ADD_ADDR allows an MPTCP host to signal a new address explicitly to the other host to allow it to initiate a new subflow - as an alternative to using MP_JOIN to directly start a new subflow). Implementation 3 decided not to support sending ADD_ADDR or processing ADD_ADDR as it is considered a security risk. Implementation 4 decided not to support ADD_ADDR because it didn't think it would be useful as most clients are behind NATing devices. However, both implemented REMOVE_ADDR (in Implementation 3 the client can send a REMOVE_ADDR but ignores incoming REMOVE_ADDR).

In Implementations 1, 3 and 4 only the initiator of the original subflow can start a new subflow (a reason mentioned is that NATs make it hard for the server to reach the client).

All implementations support 4 bytes "Data ACK" and "Data sequence number" fields, and will interoperate with an implementation sending 8 bytes. Implementation 1 uses only 4 bytes fields; if an implementation sends an 8 byte data sequence number it replies with a 4 byte data ack.

3.4. Question 4: Fallback from MPTCP

Question 4 asks about action when there is a problem with MPTCP, for example due to a middlebox mangling MPTCP's signalling. The connection needs to fall back: if the problem is on the first subflow then MPTCP falls back to TCP, whilst if the problem is on an

additional subflow then that subflow is closed with a TCP RST, as discussed in [Section 3.6 RFC6824].

Implementations 3 and 4 made several comments about fallback.

Implementation 3 suggests that both sender and receiver behaviours could be outlined with more detail, in particular when DSS checksum is not in use and the MPTCP options are stripped. Implementation 3 falls back to TCP when there's one sub flow, but not when there are multiple sub flows (MPTCP is used in active/backup mode, and it is assumed that the sub flow transferring data is most likely to be more usable than any other established sub flow, hence the sub flow on which fallback occurred is kept alive and other sub flows are closed).

Implementation 4 found a corner case where it is not clear what to do: if a pure ack or data packet without DSS is received in middle of transaction (which can happen if the routing changes and the new path drops MPTCP options). Also, Implementation 4 suggests that clarifying whether the infinite map exchange is unidirectional or bidirectional.

Implementation 1 has developed a publicly available test suite that tests MPTCP's traversal of middleboxes.

3.5. Question 5: Heuristics

Question 5 gathers information about heuristics: aspects that are not required for protocol correctness but impact the performance. We would like to document best practice so that future implementers can learn from the experience of pioneers.

There are several differences between the implementations.

For receiver buffer, Implementation 1 uses a slightly modified version of Linux's auto-tuning algorithm; Implementation 2 determines the receiver buffer by using "TCP_MAXWIN << tp->rcv_scale" (this is a temporary measure); Implementation 3 uses MPTCP in active/backup mode, so the receive buffer sizes at the MPTCP and subflow level is the same (automatic buffer tuning is turned off); Implementation 4 varies the receiver buffer size based on the services and application type.

For the sender buffer, Implementation 1 uses Linux auto-tuning, Implementation 2 scales based on occupancy, whilst Implementation 3 turns off automatic buffer tuning, and Implementation 4 uses MPTCP-level (sub)flow control that is (almost) the same as regular TCP flow control.

Implementations 1, 2 and 3 re-transmit unacknowledged data on a different subflow (and not the same subflow), whilst Implementation 4 re-transmits on original subflow for 3 RTOs and then uses another subflow.

For port usage, Implementations 1 and 2 uses the same ports for the additional subflows, whilst Implementation 3 uses the same destination port but a different local port, so that on the wire it looks like two connections to the same remote destination.

Implementation 4 suggests that the RFC should more clearly /extensively define failure cases and how to handle unexpected signals.

3.6. Question 6: Security

Question 6 asks about security related matters.

All Implementations have implemented the hash-based, HMAC-SHA1 security mechanism defined in [RFC6824]. Implementation 3 suggests that a more secure mechanism could be tied with SSL. Implementation 4 suggests that a more secure and lightweight mechanism is needed, as keys are exchanged (in the MP_CAPABLE option) in plain text and the key generation mechanism is highly computational intensive. Implementation 1 has implemented two additional mechanisms in a separate Linux branch - one lightweight and the other SSL-based.

3.7. Question 7: IANA

All Implementations have followed the IANA-related definitions [Section 8 RFC6824] for: TCP Option Kind number (30); the sub-registry for "MPTCP Option Subtypes"; and the sub-registry for "MPTCP Handshake Algorithms".

3.8. Question 8: Congestion control and subflow policy

Question 8 asks how is shared across multiple subflows.

Implementation 1 has added support for coupled congestion control (both that defined in [RFC6356] and in OLIA, draft-khalili-mptcp-congestion-control. The other implementations do not include coupled congestion control. Whilst Implementation 2 plans to add it (currently it uses a simple algorithm spreads traffic across the subflows), Implementations 3 and 4 do not plan to add coupled congestion control - they use one subflow at a time, with others as a backup. Implementation 3 believes it is not currently useful to share load across all network interfaces on a mobile node, as the interfaces have different characteristics for cost, bring-up

and power usage. They have both found the B bit (in MP-JOIN) and MP-PRIO option very useful for this active /backup operation.

Implementation 2 is also interested in experimenting with congestion control across paths with different path-cost metrics.

3.9. Question 9: API

Question 9 gathers information about the API. None have implemented the [RFC6897] "basic MPTCP API" for MPTCP-aware applications. For three implementations MPTCP is used for all applications (set by configuration), whilst Implementation 3 uses a private API that allows MPTCP to be used on a per application basis.

3.10. Question 10: Deployments, use cases and operational experiences

Question 10 takes the opportunity of this survey to gather some limited information about operational experiences and deployments.

The Implementations mention different use cases.

Implementation 2 is interested in using MPTCP for several use cases: vehicle to infrastructure (V2I) connectivity (to provide a persistent connection using 3G and roadside wifi); multi-homed "home-user" environments; high throughput data transfers. Implementation 3 is interested in the mobile scenario, with MPTCP providing an active /backup mode so achieving session continuity across changing network environments. Implementation 4 is interested in MPTCP giving reliability and fault tolerance via a proxy. Implementation 1 already uses MPTCP on www.multipath-tcp.org and for internal ssh servers at UCLouvain.

Implementation 4 uses a proxy (MPTCP connections from a client are terminated and the TCP connection established on the other side), whilst the other Implementations are on end hosts. Implementation 2 is so far within controlled testbeds, whilst Implementation 3 is on the Internet.

Implementation 2 is currently an alpha-quality build, so limited testing so far.

Implementation 3 suggests working at the IETF with firewall vendors, to get them to change their defaults to allow MPTCP signals. This would also reduce the "over-engineering" needed to handle fallback cases. Implementation 1 suggests retrieving logs from middleboxes, as the best approach to understanding the interactions of MPTCP signalling with middleboxes.

Implementation 3 discusses a scenario that should be handled better. A backup subflow may never sent data. If the initial subflow fails, data is retransmitted on the backup subflow, but that path has a middlebox stripping options. Then it may not be possible to recover the MPTCP session.

3.11. Question 11: Improvements to RFC6824

Question 11 asks if there are any areas where RFC6824 could be improved. The main topics have been mentioned earlier:

- o fallback: the need for more clarity in the fallback cases is mentioned by Implementations 3 and 4.
- o security: the need for both a more secure and a more lightweight mechanism is mentioned.

Implementation 3 also suggests several potential improvements, which are outside the scope of RFC6824: support for sub flow level automatic buffer scaling, varying QoS support, and varying window scaling support on each sub flow; also, additional work on option signalling will be brought up in future discussions.

4. IANA Considerations

This document makes no request of IANA.

5. Security Considerations

This survey does not impact the security of MPTCP, except to the extent that it uncovers security issues that can be tackled in a future version of the protocol.

6. Acknowledgements

Many thanks to the people who replied to the survey: Christoph Paasch, Nigel Williams, anon, and Krishna Khanal. Very many thanks to all of the teams who actually did the implementation and testing and are continuing to improve them.

7. Full survey response for Implementation 1

Question 1: Your details Question 1 gathers some information about the team that has implemented MPTCP.

1. Your institution: UCLouvain, IP Networking Lab
(<http://inl.info.ucl.ac.be>)

2. Name(s) of people in your implementation and test teams: Initial design from Sebastien Barre. Since then, numerous code-contributors (ordered by number of commits): Christoph Paasch (UCLouvain) Gregory Detal (UCLouvain) Jakko Korkeaniemi (Aalto University) Mihai P. Andrei (Intel) Fabien Duchene (UCLouvain) Andreas Seelinger (RWTH Aachen) Stefan Sicleru (Intel) Lavkesh Lahngir Catalin Nicutar (PUB Bucharest) Andrei Maruseac (Intel) Andreas Ripke (NEC) Vlad Dogaru (Intel) Octavian Purdila (Intel) Niels Laukens (VRT Belgium) John Ronan (TSSG) Brandon Heller (Stanford University) Conformance Testing: Yvan Coene (UCLouvain)

3. Do you want your answers to Question 1.1 and 1.2 above to be anonymised? No.

3.2. Question 2: Preliminary information about your implementation
Question 2 gathers some preliminary information.

1. What OS is your implementation for? (or is it application layer?)
Linux Kernel.

2. Do you support IPv4 or IPv6 addresses or both? We support both.

3. Is it publicly available (or will it be?) (for free or a fee?)
Publicly available (GPLv2) at www.multipath-tcp.org

4. Overall, what are you implementation and testing plans? (details can be given against individual items later) We plan to continue to align our implementation with the IETF specifications and improve its performance and stability.

5. Is it an independent implementation? Or does it build on another MPTCP implementation -which one? Independent implementation.

6. Have you already done some interop tests, for example with UCLouvain's "reference" Linux implementation? /

7. Would you be prepared to take part in an interop event, for example adjacent to IETF-87 in Berlin? Yes. We are also ready to help in organising such an event if needed.

3.3. Question 3: Support for MPTCP's Signalling Functionality
Question 3 asks about support for the various signalling messages that the MPTCP protocol defines. *** For each message, please give a little information about the status of your implementation: for example, you may have implemented it and fully tested it; the

implementation may be in progress; you have not yet implemented it but plan to soon (timescale?); you may you have no intention to implement it (why?); etc.

1. Connection initiation (MP_CAPABLE) [Section 3.1 RFC6824]

a. What is the status of your implementation? Fully support the MP_CAPABLE exchange.

b. Any other comments or information? We generate the random key as a hash of the 5-tuple, sequence number and a local secret. This significantly improves the performance, instead of using a pseudo-random number generator. The performance benefit has been shown during IETF85
<http://tools.ietf.org/agenda/85/slides/slides-85-mptcp-2.pdf>

2. Starting a new subflow (MP_JOIN) [Section 3.2 RFC6824]

a. What is the status of your implementation? Fully support the MP_JOIN exchange.

b. Can either end of the connection start a new subflow (or only the initiator of the original subflow)? Currently, only the initiator of the original subflow starts a new subflow. Given the widespread deployment of NATs, it is often difficult for the server to reach the client. This is the main reason why the server currently does not start new subflows in our implementation. But, the initiator would accept a SYN+MP_JOIN if sent by another implementation.

c. What is the maximum number of subflows your implementation can support? Currently 32.

d. Any other comments or information?

3. Data transfer (DSS) [Section 3.3 RFC6824]

a. What is the status of your implementation? Fully working implementation of data transfer.

b. The "Data ACK" field can be 4 or 8 octets. Which one(s) have you implemented? We use 4 bytes for the DATA-ACK field.

c. The "Data sequence number" field can be 4 or 8 octets. Which one(s) have you implemented? We use 4 bytes for the data sequence number.

d. Does your implementation support the "DATA_FIN" operation to close an MPTCP connection? Yes.

e. Does your implementation support the "Checksum" field (which is negotiated in the MP_CAPABLE handshake)? Yes. This is configurable via a sysctl.

f. Any other comments or information? We support interoperability with implementations that do send 64-bit data sequence numbers and data acks. However, even if the peer sends 64-bit data sequence numbers, we will only reply with a 32-bit data-ack. We do not have heuristics to trigger the sending of DATA_ACKs. We simply send the DATA_ACK in each packet.

4. Address management (ADD_ADDR and REMOVE_ADDR) [Section 3.4 RFC6824]

a. What is the status of your implementation? We support ADD_ADDR/REMOVE_ADDR messages.

b. Can your implementation do ADD_ADDRESS for addresses that appear *after* the connection has been established? Yes, as shown in: "Exploring Mobile/WiFi Handover with Multipath TCP", C. Paasch et. al, ACM SIGCOMM workshop on Cellular Networks (Cellnet'12), 2012.

c. Any other comments or information? We do not send out TCP keepalive-messages upon the reception of a REMOVE_ADDR-message.

5. Fast close (MP_FASTCLOSE) [Section 3.5 RFC6824]

a. What is the status of your implementation? We support the MP_FASTCLOSE implementation.

b. Any other comments or information?

3.4. Question 4: Fallback from MPTCP Question 4 asks about action when there is a problem with MPTCP, for example due to a middlebox mangling MPTCP's signalling. The connection needs to fall back: if the problem is on the first subflow then MPTCP falls back to TCP, whilst if the problem is on an additional subflow then that subflow is closed with a TCP RST, as discussed in [Section 3.6 RFC6824].

1. If the MP_CAPABLE option is removed by a middlebox, does your implementation fall back to TCP? Yes.

2. If the MP_JOIN option does not get through on the SYNs, does your implementation close the additional subflow? Yes.

3. If the DSS option does not get through on the first data segment(s), does your implementation fall back? (either falling back to MPTCP (if the issue is on the first subflow) or closing the

additional subflow (if the issue is on an additional subflow)) Yes. On the initial subflow we do a seamless fallback, additional subflows will be closed by a RST.

4. Similarly, if the "DATA ACK" field does not correctly acknowledge the first data segment(s), does your implementation fall back? Yes. Same as above.

5. Does your implementation protect data with the "Checksum" field in the DSS option [Section 3.3 RFC6824]? If the checksum fails (because the subflow has been affected by a middlebox), does your implementation immediately close the affected subflow (with a TCP RST) with the MP_FAIL Option? If the checksum fails and there is a single subflow, does your implementation handle this as a special case, as described in [Section 3.6 RFC6824]? Yes, we support the DSS-checksum. If the checksum is wrong and there exist other subflows, we close the current subflow with an RST. If there is no other subflow, we send an ACK + MP_FAIL and do a fallback to infinite mapping. This fallback has successfully been tested with different type of NAT middleboxes, while using FTP.

6. Does your implementation fall back to TCP by using an "infinite mapping" [Section 3.3.1 RFC6824] (so that the subflow-level data is mapped to the connection-level data for the remainder of the connection)? Yes.

7. Did you find any corner cases where MPTCP's fallback didn't happen properly? No. We have developed a test-suite to test the middlebox-traversal of MPTCP, available at <http://multipath-tcp.org/pmwiki.php/Users/AboutMeasures>

8. Any other comments or information about fallback?

3.5. Question 5: Heuristics Question 5 gathers information about heuristics: aspects that are not required for protocol correctness but impact the performance. We would like to document best practice so that future implementers can learn from the experience of pioneers. The references contain some initial comments about each topic.

1. Receiver considerations [S3.3.4, RFC6824]: What receiver buffer have you used? Does this depend on the retransmission strategy? What advice should we give about the receiver? Linux includes an autuning algorithm for the TCP receiver buffer. This algorithm has been slightly modified for Multipath TCP. The receive-buffer does not depend on the retransmission strategy.

2. Sender considerations [S3.3.5, RFC6824]: How do you determine how

much data a sender is allowed to send and how big the sender buffer is? What advice should we give about the sender? The send-buffer is autotuned similarly as the receive-buffer (see above). We send as much data as possible, filling the congestion windows of each subflow. The sender deploys the "Opportunistic Retransmission" and "Penalization" algorithms from the paper: "How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP", C. Raiciu et. al, NSDI 2012.

3. Reliability and retransmissions [S3.3.6, RFC6824]: What is your retransmission policy? (when do you retransmit on the original subflow vs on another subflow or subflows?) When do you decide that a subflow is underperforming and should be reset, and what do you then do? What advice should we give about this issue? Upon an RTO on subflow A, we reinject all the unacknowledged data of subflow A on another subflows. We do not currently have a mechanism to detect that a subflow is underperforming.

4. Port usage [S3.3.8.1, RFC6824]: Does your implementation use the same port number for additional subflows as for the first subflow? Have you used the ability to define a specific port in the Add Address option? What advice should we give about this issue? We always use the same port number as for the first subflow. Except, if the ADD_ADDRESS option that has been received contained a specific port. We do not have a means to configure the specific port in the ADD_ADDRESS option, but we support reception of the port.

5. Delayed subflow start [S3.3.8.2, RFC6824]: What factors does your implementation consider when deciding about opening additional subflows? What advice should we give about this issue? As soon as we are sure that the initial subflow is fully MPTCP-capable (reception of a DATA_ACK), we create a full mesh among all IP-addresses between the two hosts. We do not explicitly delay the creation of new subflows.

6. Failure handling [S3.3.8.3, RFC6824]: Whilst the protocol defines how to handle some unexpected signals, the behaviour after other unexpected signals is not defined. What advice should we give about this issue? We did not implement the caching mentioned in Section 3.8.3.

7. Use of TCP options: As discussed in [Appendix A, RFC6824], the TCP option space is limited, but a brief study found there was enough room to fit all the MPTCP options. However there are constraints on which MPTCP option(s) can be included in packets with other TCP options - do the suggestions in Appendix A need amending or expanding? We do not implement specific heuristics to reduce the TCP option-space usage. If timestamp is enabled we will only be able to

send two SACK-blocks, because the DATA_ACK consumes the remaining bytes.

8. What other heuristics should we give advice about? Any other comments or information?

3.6. Question 6: Security Question 6 asks about Security related matters [Section 5 RFC6824].

1. Does your implementation use the hash-based, HMAC-SHA1 security mechanism defined in [RFC6824]? Yes.

2. Does your implementation support any other handshake algorithms? We have in a separate branch, an implementation of draft-paasch-mptcp-lowoverhead and draft-paasch-mptcp-ssl.

3. It has been suggested that a Standards-track MPTCP needs a more secure mechanism. Do you have any views about how to achieve this? We believe that the solution described in draft-paasch-mptcp-ssl would be a good starting point since it leverages the security of the upper layer.

4. Any other comments or information?

3.7. Question 7: IANA Question 7 asks about IANA related matters.

1. Does your implementation follow the IANA-related definitions? [Section 8 RFC6824] defines: TCP Option Kind number (30); the sub-registry for "MPTCP Option Subtypes"; and the sub-registry for "MPTCP Handshake Algorithms" Yes.

2. Any other comments or information?

3.8. Question 8: Congestion control and subflow policy Question 8 asks about how you share traffic across multiple subflows.

1. How does your implementation share traffic over the available paths? For example: as a spare path on standby ('all-or-nothing'), as an 'overflow', etc? Does it have the ability to send /receive traffic across multiple subflows simultaneously? The implementation is able to send and receive traffic on all subflows simultaneously. Our scheduler first tries to send traffic on the subflow with the lowest RTT. As this subflow's congestion window is full, we pick the subflow with the next lower RTT.

2. Does your implementation support "handover" from one subflow to another when losing an interface? Yes, as described in: "Exploring Mobile/WiFi Handover with Multipath TCP", C. Paasch et. al, ACM

SIGCOMM workshop on Cellular Networks (Cellnet'12), 2012.

3. Does your implementation support the coupled congestion control defined in [RFC6356]? Yes.

4. Does your implementation support some other coupled congestion control (ie that balances traffic on multiple paths according to feedback)? We also support the OLIA congestion control (draft-khalili-mptcp-congestion-control-00).

5. The MP_JOIN (Starting a new subflow) Option includes the "B" bit, which allows the sender to indicate whether it wishes the new subflow to be used immediately or as a backup if other path(s) fail. The MP_PRIO Option is a request to change the "B" bit - either on the subflow on which it is sent, or (by setting the optional Address ID field) on other subflows. Does your implementation support the "B" bit and MP_PRIO mechanisms? Do you think they're useful, or have another suggestion? Yes, we support the "B"-bit of the MP_JOIN and the MP_PRIO option. It is configurable on a per-interface basis. Experiences with the "B"-bit can be found in our paper: "Exploring Mobile/WiFi Handover with Multipath TCP", C. Paasch et. al, ACM SIGCOMM workshop on Cellular Networks (Cellnet'12), 2012.

6. Any other comments or information or suggestions about the advice we should give about congestion control [S3.3.7 RFC6824] and subflow policy [S3.3.8 RFC6824]?

3.9. Question 9: API Question 9 gathers information about your API. [RFC6897] considers the MPTCP Application Interface.

1. With your implementation, can legacy applications use (the existing sockets API to use) MPTCP? How does the implementation decide whether to use MPTCP? Should the advice in [Section 4, RFC6897] be modified or expanded? Yes, a standard TCP socket API can be used. By default MPTCP is enabled on all connections.

2. The "basic MPTCP API" enables MPTCP-aware applications to interact with the MPTCP stack via five new socket options. For each one, have you implemented it? has it been useful? None of them are part of the current stable release MPTCP v0.86.
<http://multipath-tcp.org/pmwiki.php?n=Main.Release86> a.
TCP_MULTIPATH_ENABLE? b. TCP_MULTIPATH_ADD? c.
TCP_MULTIPATH_REMOVE? d. TCP_MULTIPATH_SUBFLOWS? e.
TCP_MULTIPATH_CONNID?

3. Have you implemented any aspects of an "advanced MPTCP API"? ([Appendix A, RFC6897] hints at what it might include.) No.

4. Any other comments or information?

3.10. Question 10: Deployments, use cases and operational experiences
Question 10 takes the opportunity of this survey to gather some limited information about operational experiences and deployments. Any very brief information would be appreciated, for example: 1. What deployment scenarios are you most interested in? 2. Is your deployment on "the Internet" or in a controlled environment? 3. Is your deployment on end hosts or with a MPTCP-enabled proxy (at one or both ends)? 4. What do you see as the most important benefits of MPTCP in your scenario(s)? 5. How extensively have you deployed and experimented with MPTCP so far?

Our implementation is open-source and has been discussed for various types of tests/deployments based on the messages received on the mptcp-dev mailing list. We currently use Multipath TCP on www.multipath-tcp.org and also on internal ssh servers at UCLouvain. 6. MPTCP's design seeks to maximise the chances that the signalling works through middleboxes. Did you find cases where middleboxes blocked MPTCP signalling? We have implemented a test suite based on a slightly modified version of the Multipath TCP implementation that allows to check the interoperability between Multipath TCP and middleboxes. We have used it over Internet paths and identified some potential problems. However, the best approach to test these interactions would be to control the middlebox and analyse its logs during the Multipath TCP test. The test suite can be retrieved from <http://multipath-tcp.org/pmwiki.php/Users/AboutMeasures>

7. MPTCP's design seeks to ensure that, if there is a problem with MPTCP signalling, then the connection either falls back to TCP or removes the problematic subflow. Did you find any corner cases where this didn't happen properly? See above.

8. Have you encountered any issues or drawbacks with MPTCP?

9. Any other comments or information?

3.11. Question 11: Improvements to RFC6824

1. Are there any areas where [RFC6824] could be improved, either in technical content or clarity? 2. Any other issues you want to raise?

8. Full survey response for Implementation 2

Question 1: Your details

1.1 Swinburne University of Technology, Hawthorn, Victoria, Australia

1.2 Lawrence Stewart, Nigel Williams

1.3 No

Question 2: Preliminary information about your implementation

2.1 FreeBSD-10

2.2 Currently IPv4 only (IPv6 support will eventually be added)

2.3 Publicly available (<http://caia.swin.edu.au/urp/newtcp/mptcp/>).
The code is released under the BSD license. 2.3

2.5 Independent

2.6 Yes, some limited testing to establish interoperability.

2.7 Yes, with some additional work this should be possible (if not then IETF-88). Q

Question 3: Support for MPTCP's Signaling Functionality

3.1 a) MP_CAPABLE Implemented

b) Do not currently honour checksum flag (to be implemented)

3.2 a) MP_JOIN Implemented

b) Either end can initiate a MP_JOIN

c) 8 (controlled via sysctl)

d) Currently do not include HMAC verification during handshake, but this will be enabled in the next patch (several weeks from time of submission)

3.3 a) DSS Implemented

b) 4 (default) and 8

c) 4 (default) and 8

d) Yes, however the connection tear-down exchange is not fully implemented - the connection shuts down but the DFIN may not be

correctly acknowledged.

e) No. This will be supported eventually (time-frame unknown)

3.4 a) ADD_ADDR implemented, REMOVE_ADDR not implemented (to be done, timeframe unknown)

b) No. Functionality to be added

3.5 MP_FASTCLOSE not implemented. Plan to implement eventually

Question 4: Fallback from MPTCP

4.1 Yes

4.2 The subflow PCBs remain allocated, however the subflow is not used to send data.

4.3 No, tbd

4.4 No, tbd

4.5 No, checksumming not implemented

4.6 Yes

4.8 Fallback hasn't really been put through any structured tests yet

Question 5: Heuristics

5.1 We use "TCP_MAXWIN << tp->rcv_scale". This is temporary and we will use a call into the "multipath" control layer to determine this value in future releases (we need to investigate a suitable way of calculating this).

5.2 cwnd determines the amount of data to send (given that rcv window is always very large). Sendbuffer is scaled based on occupancy.

5.3 We currently don't have Data-level retransmits enabled. However our policy is to retransmit on the next subflow that requests data to send that is suitable. There is no intelligence in the packet scheduler currently,

5.4 The same port numbers are re-used for additional subflows.

Question 6: Security

6.1 Yes

6.2 No

Question 7: IANA

7.1 Yes

Question 8: Congestion Control and subflow policy

8.1 A simple algorithm is used to divide the send buffer between subflows, so that traffic is spread across the subflows.

8.3 No. (to be added)

8.4 No

8.5 No

Question 9: API

9.1 Legacy applications are able to use MPTCP. MPTCP is set globally via a sysctl variable.

9.2 No

9.3 No

Question 10: API

10.1 Some current project work is based on MPTCPs use in vehicle to infrastructure (V2I) connectivity (to provide a persistent connection using 3G and roadside wifi). Other interests are in multi-homed "home-user" environments, high throughput data transfers.... We are also interested in experimenting with congestion control across paths with different path-cost metrics.

10.2 So far only within controlled testbeds

10.3 End hosts

10.4 Depending on the scenario, connection persistence, throughput...

10.5 Still an alpha-quality build, so limited testing so far.

9. Full survey response for Implementation 3

Survey 3.1. Question 1: Your details Question 1 gathers some information about the team that has implemented MPTCP.

1. Your institution: anonymized.

2. Name(s) of people in your implementation and test teams: There were several folks involved in the implementation and testing.

3. Do you want your answers to Question 1.1 and 1.2 above to be anonymised? Yes.

3.2. Question 2: Preliminary information about your implementation Question 2 gathers some preliminary information.

1. What OS is your implementation for? (or is it application layer?) anonymized (commercial OS)

2. Do you support IPv4 or IPv6 addresses or both? Both.

3. Is it publicly available (or will it be?) (for free or a fee?) No.

4. Overall, what are you implementation and testing plans? (details can be given against individual items later) We plan to use it in a mobile environment.

5. Is it an independent implementation? Or does it build on another MPTCP implementation -which one? It is an independent implementation.

6. Have you already done some interop tests, for example with UCLouvain's "reference" Linux implementation? Most MPTCP option formats were tested with the reference Linux implementation.

7. Would you be prepared to take part in an interop event, for example adjacent to IETF-87 in Berlin? Unsure at this point.

3.3. Question 3: Support for MPTCP's Signalling Functionality Question 3 asks about support for the various signalling messages that the MPTCP protocol defines. *** For each message, please give a little information about the status of your implementation: for example, you may have implemented it and fully tested it; the implementation may be in progress; you have not yet implemented it

but plan to soon (timescale?); you may you have no intention to implement it (why?); etc.

1. Connection initiation (MP_CAPABLE) [Section 3.1 RFC6824] a. What is the status of your implementation? Fully implemented and tested against the reference Linux implementation.

b. Any other comments or information?

2. Starting a new subflow (MP_JOIN) [Section 3.2 RFC6824] a. What is the status of your implementation? Fully implemented and tested against the reference Linux implementation.

b. Can either end of the connection start a new subflow (or only the initiator of the original subflow)? Only the initiator of the original sub flow can start other sub flows.

c. What is the maximum number of subflows your implementation can support? There is no hard limit.

d. Any other comments or information?

3. Data transfer (DSS) [Section 3.3 RFC6824] a. What is the status of your implementation? Fully implemented and tested.

b. The "Data ACK" field can be 4 or 8 octets. Which one(s) have you implemented? Both have been implemented but the use of the 4-byte field is the default. When an 8 byte DSS is received, an 8 byte Data ACK is sent in response.

c. The "Data sequence number" field can be 4 or 8 octets. Which one(s) have you implemented? Both have been implemented but the use of the 4-byte field is the default. When a wraparound of the lower 32-bit part of the DSS is detected, the full 8 byte DSS is sent.

d. Does your implementation support the "DATA_FIN" operation to close an MPTCP connection? Yes. There are cases however where the sub flows are closed (TCP FIN'd) but the DATA_FIN is not sent - in this case the MPTCP connection must be closed through a garbage collector after some idle time.

e. Does your implementation support the "Checksum" field (which is negotiated in the MP_CAPABLE handshake)? Yes.

f. Any other comments or information?

4. Address management (ADD_ADDR and REMOVE_ADDR) a. What is the status of your implementation? It does not support sending ADD_ADDR

or processing ADD_ADDR as it is considered a security risk. Also, we only have a client side implementation at the moment which always initiates the sub flows. The remote end does not send ADD_ADDR in our configuration. The client can send REMOVE_ADDR however when one of the established sub flow's source address goes away. The client ignores incoming REMOVE_ADDR options also.

b. Can your implementation do ADD_ADDRESS for addresses that appear *after* the connection has been established? No. c. Any other comments or information?

5. Fast close (MP_FASTCLOSE) [Section 3.5 RFC6824] a. What is the status of your implementation? It is supported. Though Retransmission of Fast close is not supported yet.

b. Any other comments or information?

3.4. Question 4: Fallback from MPTCP Question 4 asks about action when there is a problem with MPTCP, for example due to a middlebox mangling MPTCP's signalling. The connection needs to fall back: if the problem is on the first subflow then MPTCP falls back to TCP, whilst if the problem is on an additional subflow then that subflow is closed with a TCP RST, as discussed in [Section 3.6 RFC6824].

1. If the MP_CAPABLE option is removed by a middlebox, does your implementation fall back to TCP? Yes.

2. If the MP_JOIN option does not get through on the SYNs, does your implementation close the additional subflow? Yes.

3. If the DSS option does not get through on the first data segment(s), does your implementation fall back? (either falling back to MPTCP (if the issue is on the first subflow) or closing the additional subflow (if the issue is on an additional subflow)) Yes it falls back to TCP when there's one sub flow. When there are multiple sub flows, since MPTCP is used in active/backup mode, it is assumed that the sub flow transferring data is most likely to be more usable than any other established sub flow. So the sub flow on which fallback occurred is kept alive and other sub flows are closed. Fallback though is not guaranteed to occur safely when there are more than one sub flows because the infinite mapping option may be stripped like other DSS options and the MP_FAIL option if used in scenarios other than for reporting checksum failure can also be stripped.

4. Similarly, if the "DATA ACK" field does not correctly acknowledge the first data segment(s), does your implementation fall back? No. Current implementation just ignores the unexpected data ack.

5. Does your implementation protect data with the "Checksum" field in the DSS option [Section 3.3 RFC6824]? If the checksum fails (because the subflow has been affected by a middlebox), does your implementation immediately close the affected subflow (with a TCP RST) with the MP_FAIL Option? If the checksum fails and there is a single subflow, does your implementation handle this as a special case, as described in [Section 3.6 RFC6824]? Yes.

6. Does your implementation fall back to TCP by using an "infinite mapping" [Section 3.3.1 RFC6824] (so that the subflow-level data is mapped to the connection-level data for the remainder of the connection)? Yes.

7. Did you find any corner cases where MPTCP's fallback didn't happen properly? If the very first sub flow does not send any data and is disconnected right away, then the current implementation allows a join to occur with the addition of another sub flow which then becomes a fully mp capable sub flow. Thus we allow break before make by letting additional sub flows to be joined if the very first one disconnected even without sending any data. This is a very corner case but an instance where we do not follow the rules of fallback (allow second sub flow even when first sub flow did not send/receive data/data acks).

8. Any other comments or information about fallback? Fallback after connection establishment and after a few data packets were transferred with MPTCP options is complicated. The spec does not clearly cover the cases of options being stripped by middle boxes. It goes into good detail about what to do when the DSS checksum fails, but not when DSS checksum is not in use and the MPTCP options are stripped. Both sender/receiver behaviors could be outlined with more detail.

3.5. Question 5: Heuristics Question 5 gathers information about heuristics: aspects that are not required for protocol correctness but impact the performance. We would like to document best practice so that future implementers can learn from the experience of pioneers. The references contain some initial comments about each topic.

1. Receiver considerations [S3.3.4, RFC6824]: What receiver buffer have you used? Does this depend on the retransmission strategy? What advice should we give about the receiver? We are just using MPTCP in active/backup mode. This mode is simpler wrt receive buffer utilization. The receive buffer sizes at the MPTCP and sub flow level is the same. Automatic buffer tuning is turned off when MPTCP is in use.

2. Sender considerations [S3.3.5, RFC6824]: How do you determine how much data a sender is allowed to send and how big the sender buffer is? What advice should we give about the sender? Automatic buffer tuning is turned off when MPTCP is in use.

3. Reliability and retransmissions [S3.3.6, RFC6824]: What is your retransmission policy? (when do you retransmit on the original subflow vs on another subflow or subflows?) When do you decide that a subflow is underperforming and should be reset, and what do you then do? What advice should we give about this issue? Retransmissions at MPTCP level do not occur on the same sub flow except when MP_FAIL option is received. A sub flow is said to be underperforming when its network connectivity goes away.

4. Port usage [S3.3.8.1, RFC6824]: Does your implementation use the same port number for additional subflows as for the first subflow? Have you used the ability to define a specific port in the Add Address option? What advice should we give about this issue? The destination port is the same. The local port changes for additional sub flows so on the wire it is like two tcp connections to the same remote destination. We have not used Add Address option at all.

5. Delayed subflow start [S3.3.8.2, RFC6824]: What factors does your implementation consider when deciding about opening additional subflows? What advice should we give about this issue? The client implementation is aware of network interfaces coming up or going down and establishes new sub flows or removes existing sub flows accordingly.

6. Failure handling [S3.3.8.3, RFC6824]: Whilst the protocol defines how to handle some unexpected signals, the behaviour after other unexpected signals is not defined. What advice should we give about this issue? Fallback, post establishment is probably a case that needs to be more clearly defined.

7. Use of TCP options: As discussed in [Appendix A, RFC6824], the TCP option space is limited, but a brief study found there was enough room to fit all the MPTCP options. However there are constraints on which MPTCP option(s) can be included in packets with other TCP options - do the suggestions in Appendix A need amending or expanding? Looks good already.

8. What other heuristics should we give advice about? Any other comments or information?

3.6. Question 6: Security Question 6 asks about Security related matters [Section 5 RFC6824].

1. Does your implementation use the hash-based, HMACSHA1 security mechanism defined in [RFC6824]? Yes.
2. Does your implementation support any other handshake algorithms? No.
3. It has been suggested that a Standards-track MPTCP needs a more secure mechanism. Do you have any views about how to achieve this? No. But the mechanism could be tied with SSL because SSL is used wherever security is deemed important.
4. Any other comments or information?

3.7. Question 7: IANA Question 7 asks about IANA related matters.

1. Does your implementation follow the IANA-related definitions? [Section 8 RFC6824] defines: TCP Option Kind number (30); the sub-registry for "MPTCP Option Subtypes"; and the Page 12 of 17 Survey 6/22/13, 5:55 PM sub-registry for "MPTCP Handshake Algorithms" Yes.
2. Any other comments or information? No.

3.8. Question 8: Congestion control and subflow policy Question 8 asks about how you share traffic across multiple subflows.

1. How does your implementation share traffic over the available paths? For example: as a spare path on standby ('all-or-nothing'), as an 'overflow', etc? Does it have the ability to send /receive traffic across multiple subflows simultaneously? It uses active/backup where one sub flow is preferred or has higher priority over other sub flows. When the preferred sub flow fails or begins to experience retransmission timeouts, the other sub flows are used.
2. Does your implementation support "handover" from one subflow to another when losing an interface? Yes.
3. Does your implementation support the coupled congestion control defined in [RFC6356]? No.
4. Does your implementation support some other coupled congestion control (ie that balances traffic on multiple paths according to feedback)? No.
5. The MP_JOIN (Starting a new subflow) Option includes the "B" bit, which allows the sender to indicate whether it wishes the new subflow to be used immediately or as a backup if other path(s) fail. The MP_PRIO Option is a request to change the "B" bit - either on the subflow on which it is sent, or (by setting the optional Address ID

field) on other subflows. Does your implementation support the "B" bit and MP_PRIO mechanisms? Do you think they're useful, or have another suggestion? Yes the implementation uses the B bit and the MP_PRIO option. They are very useful for the active/backup mode of operation.

6. Any other comments or information or suggestions about the advice we should give about congestion control [S3.3.7 RFC6824] and subflow policy [S3.3.8 RFC6824]?

3.9. Question 9: API Question 9 gathers information about your API. [RFC6897] considers the MPTCP Application Interface.

1. With your implementation, can legacy applications use (the existing sockets API to use) MPTCP? How does the implementation decide whether to use MPTCP? Should the advice in [Section 4, RFC6897] be modified or expanded? The implementation does not support MPTCP with existing sockets API. MPTCP is exposed through a private SPI today. If MPTCP becomes prolific over the next few years, MPTCP use shall be expanded.

2. The "basic MPTCP API" enables MPTCP-aware applications to interact with the MPTCP stack via five new socket options. For each one, have you implemented it? has it been useful? a. TCP_MULTIPATH_ENABLE? b. TCP_MULTIPATH_ADD? c. TCP_MULTIPATH_REMOVE? d. TCP_MULTIPATH_SUBFLOWS? e. TCP_MULTIPATH_CONNID? This mode of API is not used. Proprietary methods are used for achieving these basic operations.

3. Have you implemented any aspects of an "advanced MPTCP API"? ([Appendix A, RFC6897] hints at what it might include.) No.

4. Any other comments or information?

3.10. Question 10: Deployments, use cases and operational experiences Question 10 takes the opportunity of this survey to gather some limited information about operational experiences and deployments. Any very brief information would be appreciated, for example:

1. What deployment scenarios are you most interested in? MPTCP in mobile environments is very powerful when used in the active/backup mode. Since the network interfaces available on mobile devices have different cost characteristics as well as different bring up and power usage characteristics, it is not useful to share load across all available network interfaces - at least not currently. Providing session continuity across changing network environments is the key deployment scenario.

2. Is your deployment on "the Internet" or in a controlled environment? The deployment is on the Internet.
 3. Is your deployment on end hosts or with a MPTCP-enabled proxy (at one or both ends)? The deployment supports MPTCP on both ends.
 4. What do you see as the most important benefits of MPTCP in your scenario(s)? Described in point 1 of this section.
 5. How extensively have you deployed and experimented with MPTCP so far? Deployment is still in early stages. We have been experimenting with MPTCP for about a year.
 6. MPTCP's design seeks to maximise the chances that the signalling works through middleboxes. Did you find cases where middleboxes blocked MPTCP signalling? Corporate firewalls block MPTCP signaling by default. IETF is one venue where Cisco, and other firewall vendors can be asked to change their defaults to allow MPTCP signals.
 7. MPTCP's design seeks to ensure that, if there is a problem with MPTCP signalling, then the connection either falls back to TCP or removes the problematic subflow. Did you find any corner cases where this didn't happen properly? This has been covered a bit in the Fallback section. When using two sub flows in active/backup mode, there is a possibility that a backup sub flow that never sent data starts being used for retransmitting data that is not going through on the active path. While it is preferable to keep the initial sub flow that successfully sent MPTCP options and drop the backup path, the initial sub flow may be the failing one, and we may want to move to the backup path. But the backup path can be retransmitting data that did not get sent successfully on the active path and if there is a middle box in the backup sub flow's path stripping options, then we have a case where the MPTCP session may not be recoverable as it may not be evident from what point in the MPTCP sequence space, data was being sent. The spec does talk of retaining the initial sub flow and closing the failed flow. So perhaps doing the reverse is not recommended, however, it would certainly be advantageous to support MPTCP better in such a failing environment. Also, in parallel working with firewall vendors to allow MPTCP options always to not have to over-engineer these cases.
 8. Have you encountered any issues or drawbacks with MPTCP?
 9. Any other comments or information?
- 3.11. Question 11: Improvements to RFC6824 1. Are there any areas where [RFC6824] could be improved, either in technical content or clarity? Discussed in the fallback section. Other areas around

MPTCP performance such as support for sub flow level automatic buffer scaling, varying QoS support, varying window scaling support on each sub flow may be worth discussing further, although they are outside the scope of the current spec.

2. Any other issues you want to raise? Some additional work on option signaling that we will bring up in future discussions.

10. Full survey response for Implementation 4

1. Your institution: Citrix Systems, Inc.

2. Name(s) of people in your implementation and test teams: NA

3. Do you want your answers to Question 1.1 and 1.2 above to be anonymised? No

3.2. Question 2: Preliminary information about your implementation
Question 2 gathers some preliminary information.

1. What OS is your implementation for? (or is it application layer?)
NetScaler Firmware

2. Do you support IPv4 or IPv6 addresses or both? Both

3. Is it publicly available (or will it be?) (for free or a fee?)
It is available for purchase

4. Overall, what are you implementation and testing plans? (details can be given against individual items later)

5. Is it an independent implementation? Or does it build on another MPTCP implementation -which one? It is an independent implementation

6. Have you already done some interop tests, for example with UCLouvain's "reference" Linux implementation? Yes, our implementation is extensively tested with Linux reference implementation

7. Would you be prepared to take part in an interop event, for example adjacent to IETF-87 in Berlin?

3.3. Question 3: Support for MPTCP's Signalling Functionality
Question 3 asks about support for the various signalling messages that the MPTCP protocol defines. *** For each message, please give a little information about the status of your implementation: for example, you may have implemented it and fully tested it; the

implementation may be in progress; you have not yet implemented it but plan to soon (timescale?); you may you have no intention to implement it (why?); etc.

1. Connection initiation (MP_CAPABLE) [Section 3.1 RFC6824] a. What is the status of your implementation? Fully implemented and tested

b. Any other comments or information? One security concern here is that the keys are exchanged in plain text which is prone to attacks and also the key generation mechanism is highly computational intensive

2. Starting a new subflow (MP_JOIN) [Section 3.2 RFC6824] a. What is the status of your implementation? Fully implemented and tested

b. Can either end of the connection start a new subflow (or only the initiator of the original subflow)? Only the initiator of the original subflow can initiate additional subflows.

c. What is the maximum number of subflows your implementation can support? we support maximum 6 subflows.

d. Any other comments or information?

3. Data transfer (DSS) [Section 3.3 RFC6824] a. What is the status of your implementation? Fully implemented and tested

b. The "Data ACK" field can be 4 or 8 octets. Which one(s) have you implemented? Our implementation supports both 4 or 8 Octets Data Ack in both the directions

c. The "Data sequence number" field can be 4 or 8 octets. Which one(s) have you implemented? Our implementation supports both 4 or 8 Octets DSN in both the directions

d. Does your implementation support the "DATA_FIN" operation to close an MPTCP connection? YES

e. Does your implementation support the "Checksum" field (which is negotiated in the MP_CAPABLE handshake)? YES

f. Any other comments or information?

4. Address management (ADD_ADDR and REMOVE_ADDR) [Section 3.4 RFC6824]

a. What is the status of your implementation? REMOVE_ADDR is implemented and tested

b. Can your implementation do ADD_ADDRESS for addresses that appear *after* the connection has been established? NO

c. Any other comments or information? ADD_ADDRESS may not be much useful in the real environment situation given that most of the clients are behind the NATing devices.

5. Fast close (MP_FASTCLOSE) [Section 3.5 RFC6824] a. What is the status of your implementation? Implemented and tested b. Any other comments or information?

3.4. Question 4: Fallback from MPTCP Question 4 asks about action when there is a problem with MPTCP, for example due to a middlebox mangling MPTCP's signalling. The connection needs to fall back: if the problem is on the first subflow then MPTCP falls back to TCP, whilst if the problem is on an additional subflow then that subflow is closed with a TCP RST, as discussed in [Section 3.6 RFC6824].

1. If the MP_CAPABLE option is removed by a middlebox, does your implementation fall back to TCP? YES

2. If the MP_JOIN option does not get through on the SYNs, does your implementation close the additional subflow? YES

3. If the DSS option does not get through on the first data segment(s), does your implementation fall back? (either falling back to MPTCP (if the issue is on the first subflow) or closing the additional subflow (if the issue is on an additional subflow)) YES

4. Similarly, if the "DATA ACK" field does not correctly acknowledge the first data segment(s), does your implementation fall back? If the sender receives pure ack for its first DSS packet then it fallbacks to regular TCP.

5. Does your implementation protect data with the "Checksum" field in the DSS option [Section 3.3 RFC6824]? If the checksum fails (because the subflow has been affected by a middlebox), does your implementation immediately close the affected subflow (with a TCP RST) with the MP_FAIL Option? If the checksum fails and there is a single subflow, does your implementation handle this as a special case, as described in [Section 3.6 RFC6824]? Yes, our implementation supports DSS checksum and will close the subflow with RST if the checksum validation fails and there are more than one subflows and sends MP_FAIL if there is a single subflow expecting infinite map from the peer.

6. Does your implementation fall back to TCP by using an "infinite mapping" [Section 3.3.1 RFC6824] (so that the subflow-level data is

mapped to the connection-level data for the remainder of the connection)? YES.

7. Did you find any corner cases where MPTCP's fallback didn't happen properly? We have found few cases where the draft is not clear about the recommended action and fallback strategy, like: 1. what is the expected behavior when pure ack or data packet without dss is received in middle of transaction? How the hosts should fallback in this case? This can happen if the routing changes and the new path drops mptcp options. In this case MP_FAIL/infinite map exchange may not be possible and so could not decide whether both parties are in sync to fallback to tcp. 2. whether infinite map is unidirectional or bidirectional? If one host is sending infinite map to peer, does the peer also needs to send infinite map to the host? Exchanging infinite map and falling back to TCP from both ends is easy from implementation point of view. 8. Any other comments or information about fallback?

3.5. Question 5: Heuristics Question 5 gathers information about heuristics: aspects that are not required for protocol correctness but impact the performance. We would like to document best practice so that future implementers can learn from the experience of pioneers. The references contain some initial comments about each topic.

1. Receiver considerations [S3.3.4, RFC6824]: What receiver buffer have you used? Does this depend on the retransmission strategy? What advice should we give about the receiver? Our implementation uses varying buffer size based on the services and application type.

2. Sender considerations [S3.3.5, RFC6824]: How do you determine how much data a sender is allowed to send and how big the sender buffer is? What advice should we give about the sender? The send side flow control is handled at mptcp level and is independent to subflows. The mptcp level flow control is (almost) same as the regular TCP flow control.

3. Reliability and retransmissions [S3.3.6, RFC6824]: What is your retransmission policy? (when do you retransmit on the original subflow vs on another subflow or subflows?) When do you decide that a subflow is underperforming and should be reset, and what do you then do? What advice should we give about this issue? The retransmission is done by the subflows as long as the subflow is alive and is not removed by the REM_ADDR/RST/... . If 3 RTO happens on the subflow doing retransmission and multiple subflows are available then the mptcp starts retransmission from additional subflow. The original subflow continues retransmission for 7RTO and will be closed after that with RST.

4. Port usage [S3.3.8.1, RFC6824]: Does your implementation use the same port number for additional subflows as for the first subflow? Have you used the ability to define a specific port in the Add Address option? What advice should we give about this issue? Our current implementation doesnot support ADD_ADDR and subflow initiation.

5. Delayed subflow start [S3.3.8.2, RFC6824]: What factors does your implementation consider when deciding about opening additional subflows? What advice should we give about this issue? NA

6. Failure handling [S3.3.8.3, RFC6824]: Whilst the protocol defines how to handle some unexpected signals, the behaviour after other unexpected signals is not defined. What advice should we give about this issue? RFC should clearly define failure case handling otherwise it creates interoperability problems among various implementations. Our strategy in most of the unexpected failuire case is to send MP_FAIL RST with expected DSN if there are multiple subflows and MP_FAIL if there is a single subflow expecting infinite map from the peer.

7. Use of TCP options: As discussed in [Appendix A, RFC6824], the TCP option space is limited, but a brief study found there was enough room to fit all the MPTCP options. However there are constraints on which MPTCP option(s) can be included in packets with other TCP options - do the suggestions in Appendix A need amending or expanding? Looks fine now. Atleast timestamp can be included with every dss packet (28bytes for dss and 12bytes for Timestamp), but if there are any other options which needs to be included in data packets then the implementation has to choose which one to include among them.

8. What other heuristics should we give advice about? Any other comments or information?

3.6. Question 6: Security Question 6 asks about Security related matters [Section 5 RFC6824].

1. Does your implementation use the hash-based, HMAC-SHA1 security mechanism defined in [RFC6824]? YES.

2. Does your implementation support any other handshake algorithms? NO.

3. It has been suggested that a Standards-track MPTCP needs a more secure mechanism. Do you have any views about how to achieve this? Yes we also feel more secure and light weight mechanism is required.

4. Any other comments or information?

3.7. Question 7: IANA Question 7 asks about IANA related matters.

1. Does your implementation follow the IANA-related definitions? [Section 8 RFC6824] defines: TCP Option Kind number (30); the sub-registry for "MPTCP Option Subtypes"; and the sub-registry for "MPTCP Handshake Algorithms" YES. 2. Any other comments or information?

3.8. Question 8: Congestion control and subflow policy Question 8 asks about how you share traffic across multiple subflows.

1. How does your implementation share traffic over the available paths? For example: as a spare path on standby ('all-or-nothing'), as an 'overflow', etc? Does it have the ability to send /receive traffic across multiple subflows simultaneously? We give preference to the path that client is currently using to send data/ack and also has policy based on primary/backup setup. We accept data from multiple subflows simultaneously but don't send it simultaneously out.

2. Does your implementation support "handover" from one subflow to another when losing an interface? YES.

3. Does your implementation support the coupled congestion control defined in [RFC6356]? NO.

4. Does your implementation support some other coupled congestion control (ie that balances traffic on multiple paths according to feedback)? NO.

5. The MP_JOIN (Starting a new subflow) Option includes the "B" bit, which allows the sender to indicate whether it wishes the new subflow to be used immediately or as a backup if other path(s) fail. The MP_PRIO Option is a request to change the "B" bit - either on the subflow on which it is sent, or (by setting the optional Address ID field) on other subflows. Does your implementation support the "B" bit and MP_PRIO mechanisms? Do you think they're useful, or have another suggestion? YES, our implementation supports both 'B' flag and MP_PRIO options, they are much useful to change the priority of the subflows and to decide which subflow to use for data transfer.

6. Any other comments or information or suggestions about the advice we should give about congestion control [S3.3.7 RFC6824] and subflow policy [S3.3.8 RFC6824]?

3.9. Question 9: API Question 9 gathers information about your API. [RFC6897] considers the MPTCP Application Interface.

1. With your implementation, can legacy applications use (the existing sockets API to use) MPTCP? How does the implementation decide whether to use MPTCP? Should the advice in [Section 4, RFC6897] be modified or expanded? NA.
2. The "basic MPTCP API" enables MPTCP-aware applications to interact with the MPTCP stack via five new socket options. For each one, have you implemented it? has it been useful? a. TCP_MULTIPATH_ENABLE? b. TCP_MULTIPATH_ADD? c. TCP_MULTIPATH_REMOVE? d. TCP_MULTIPATH_SUBFLOWS? e. TCP_MULTIPATH_CONNID? NA.
3. Have you implemented any aspects of an "advanced MPTCP API"? ([Appendix A, RFC6897] hints at what it might include.) NA. 4. Any other comments or information?

3.10. Question 10: Deployments, use cases and operational experiences
Question 10 takes the opportunity of this survey to gather some limited information about operational experiences and deployments. Any very brief information would be appreciated, for example:

1. What deployment scenarios are you most interested in? MPTCP Proxy deployment where the mptcp connections from the clients are terminated and the tcp connection is established on the other side.
2. Is your deployment on "the Internet" or in a controlled environment? Targeted for the Internet deployment.
3. Is your deployment on end hosts or with a MPTCP-enabled proxy (at one or both ends)? Proxy.
4. What do you see as the most important benefits of MPTCP in your scenario(s)? Reliability and fault tolerance.
5. How extensively have you deployed and experimented with MPTCP so far?
6. MPTCP's design seeks to maximise the chances that the signalling works through middleboxes. Did you find cases where middleboxes blocked MPTCP signalling? Yes some firewalls seem dropping MPTCP options.
7. MPTCP's design seeks to ensure that, if there is a problem with MPTCP signalling, then the connection either falls back to TCP or removes the problematic subflow. Did you find any corner cases where this didn't happen properly? Few cases listed above.

8. Have you encountered any issues or drawbacks with MPTCP? 9. Any other comments or information?

3.11. Question 11: Improvements to RFC6824

1. Are there any areas where [RFC6824] could be improved, either in technical content or clarity? More clarity required in fallback cases.

2. Any other issues you want to raise?

11. Normative References

[RFC6356] Raiciu, C., Handley, M., and D. Wischik, "Coupled Congestion Control for Multipath Transport Protocols", RFC 6356, October 2011.

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

Author's Address

Philip Eardley
BT

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 28, 2015

M. Bagnulo
UC3M
C. Paasch
UCLouvain
F. Gont
SI6 Networks / UTN-FRH
O. Bonaventure
UCLouvain
C. Raiciu
UPB
March 27, 2015

Analysis of MPTCP residual threats and possible fixes
draft-ietf-mptcp-attacks-04

Abstract

This document performs an analysis of the residual threats for MPTCP and explores possible solutions to them.

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 28, 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

1. Introduction	2
2. ADD_ADDR attack	4
2.1. Possible security enhancements to prevent this attack . .	10
3. DoS attack on MP_JOIN	10
3.1. Possible security enhancements to prevent this attack . .	11
4. SYN flooding amplification	11
4.1. Possible security enhancements to prevent this attack . .	12
5. Eavesdropper in the initial handshake	12
5.1. Possible security enhancements to prevent this attack . .	13
6. SYN/JOIN attack	13
6.1. Possible security enhancements to prevent this attack . .	14
7. Reccomendations	14
7.1. Security enhancements for MPTCP	15
8. Security considerations	15
9. IANA Considerations	15
10. Acknowledgments	15
11. References	15
11.1. Normative References	16
11.2. Informative References	16
Authors' Addresses	17

1. Introduction

This document provides a complement to the threat analysis for Multipath TCP (MPTCP) [RFC6824] documented in RFC 6181 [RFC6181]. RFC 6181 provided a threat analysis for the general solution space of extending TCP to operate with multiple IP addresses per connection. Its main goal was to leverage previous experience acquired during the design of other multi-address protocols, notably SHIM6 [RFC5533], SCTP [RFC4960] and MIPv6 [RFC6275] for designing MPTCP. Thus, RFC 6181 was produced before the actual MPTCP specification (RFC6824) was completed, and documented a set of reccomendations that were considered during the production of such specification.

This document complements RFC 6181 with a vulnerability analysis of the specific mechanisms specified in RFC 6824. The motivation for this analysis is to identify possible security issues with MPTCP as currently specified and propose security enhancements to address the identified security issues.

The goal of the security mechanisms defined in RFC 6824 were to make MPTCP no worse than currently available single-path TCP. We believe that this goal is still valid, so we will perform our analysis on the same grounds. This document describes all the threats identified that are specific to MPTCP (as defined in RFC6824) that are not possible with (single-path) TCP. This means that threats that are common to TCP and MPTCP are not covered in this document.

Types of attackers: for all attacks considered in this document, we identify the type of attacker. We can classify the attackers based on their location as follows:

- o Off-path attacker. This is an attacker that does not need to be located in any of the paths of the MPTCP session at any point in time during the lifetime of the MPTCP session. This means that the Off-path attacker cannot eavesdrop any of the packets of the MPTCP session.
- o Partial time On-path attacker. This is an attacker that needs to be in at least one of the paths during part but not during the entire lifetime of the MPTCP session. The attacker can be in the forward and/or backward directions, for the initial subflow and/or other subflows. The specific needs of the attacker will be made explicit in the attack description.
- o On-path attacker. This attacker needs to be on at least one of the paths during the whole duration of the MPTCP session. The attacker can be in the forward and/or backward directions, for the initial subflow and/or other subflows. The specific needs of the attacker will be made explicit in the attack description.

We can also classify the attackers based on their actions as follows:

- o Eavesdropper. The attacker is able to capture some of the packets of the MPTCP session to perform the attack, but it is not capable of changing, discarding or delaying any packet of the MPTCP session. The attacker can be in the forward and/or backward directions, for the initial subflow and/or other subflows. The specific needs of the attacker will be made explicit in the attack description.
- o Active attacker. The attacker is able to change, discard or delay some of the packets of the MPTCP session. The attacker can be in the forward and/or backward directions, for the initial subflow and/or other subflows. The specific needs of the attacker will be made explicit in the attack description.

In this document, we consider the following possible combinations of attackers:

- o an On-path eavesdropper
- o an On-path active attacker
- o an Off-path active attacker
- o a Partial-time On-path eavesdropper
- o a Partial-time On-path active attacker

In the rest of the document we describe different attacks that are possible against the MPTCP protocol specified in RFC6824 and we propose possible security enhancements to address them.

2. ADD_ADDR attack

Summary of the attack:

Type of attack: MPTCP session hijack enabling Man-in-the-Middle.

Type of attacker: Off-path, active attacker.

Description:

In this attack, the attacker uses the ADD_ADDR option defined in RFC6824 to hijack an ongoing MPTCP session and enables himself to perform a Man-in-the-Middle attack on the MPTCP session.

Consider the following scenario. Host A with address IPA has one MPTCP session with Host B with address IPB. The MPTCP subflow between IPA and IPB is using port PA on host A and port PB on host B. The tokens for the MPTCP session are TA and TB for Host A and Host B respectively. Host C is the attacker. It owns address IPC. The attack is executed as follows:

1. Host C sends a forged packet with source address IPA, destination address IPB, source port PA and destination port PB. The packet has the ACK flag set. The TCP sequence number for the segment is i and the ACK sequence number is j . We will assume all these are valid, we discuss what the attacker needs to figure these ones later on. The packet contains the ADD_ADDR option. The ADD_ADDR option announces IPC as an alternative address for the connection. It also contains an eight bit address identifier which does not bring any strong security benefit.

2. Host B receives the ADD_ADDR message and it replies by sending a TCP SYN packet. (Note: the MPTCP specification states that the host receiving the ADD_ADDR option may initiate a new subflow. If the host is configured so that it does not initiate a new subflow the attack will not succeed. For example, on the current Linux implementation, the server does not create subflows. Only the client does so.) The source address for the packet is IPB, the destination address for the packet is IPC, the source port is PB' and the destination port is PA' (It is not required that PA=PA' nor that PB=PB'). The sequence number for this packet is the new initial sequence number for this subflow. The ACK sequence number is not relevant as the ACK flag is not set. The packet carries an MP_JOIN option and it carries the token TA. It also carries a random nonce generated by Host B called RB.
3. Host C receives the SYN+MP_JOIN packet from Host B, and it alters it in the following way. It changes the source address to IPC and the destination address to IPA. It sends the modified packet to Host A, impersonating Host B.
4. Host A receives the SYN+MP_JOIN message and it replies with a SYN/ACK+MP_JOIN message. The packet has source address IPA and destination address IPC, as well as all the other needed parameters. In particular, Host A computes a valid HMAC and places it in the MP_JOIN option.
5. Host C receives the SYN/ACK+MP_JOIN message and it changes the source address to IPC and the destination address to IPB. It sends the modified packet to IPB impersonating Host A.
6. Host B receives the SYN/ACK+MP_JOIN message. Host B verifies the HMAC of the MP_JOIN option and confirms its validity. It replies with an ACK+MP_JOIN packet. The packet has source address IPB and destination address IPC, as well as all the other needed parameters. The returned MP_JOIN option contains a valid HMAC computed by Host B.
7. Host C receives the ACK+MP_JOIN message from B and it alters it in the following way. It changes the source address to IPC and the destination address to IPA. It sends the modified packet to Host A impersonating Host B.
8. Host A receives the ACK+MP_JOIN message and creates the new subflow.

At this point the attacker has managed to place itself as a MitM for one subflow for the existing MPTCP session. It should be noted that there still exists the subflow between

address IPA and IPB that does not flow through the attacker, so the attacker has not completely intercepted all the packets in the communication (yet). If the attacker wishes to completely intercept the MPTCP session it can do the following additional step.

9. Host C sends two TCP RST messages. One TCP RST packet is sent to Host B, with source address IPA and destination address IPB and source and destination ports PA and PB, respectively. The other TCP RST message is sent to Host A, with source address IPB and destination address IPA and source and destination ports PB and PA, respectively. Both RST messages must contain a valid sequence number. Note that figuring the sequence numbers to be used here for subflow A is the same difficulty as being able to send the initial ADD_ADDR option with valid Sequence number and ACK value. If there are more subflows, then the attacker needs to find the Sequence Number and ACK for each subflow.

At this point the attacker has managed to fully hijack the MPTCP session.

Information required by the attacker to perform the described attack:

In order to perform this attack the attacker needs to guess or know the following pieces of information: (The attacker need this information for one of the subflows belonging to the MPTCP session.)

- o the four-tuple {Client-side IP Address, Client-side Port, Server-side Address, Servcer-side Port} that identifies the target TCP connection
- o a valid sequence number for the subflow
- o a valid ACK sequence number for the subflow
- o a valid address identifier for IPC

TCP connections are uniquely identified by the four-tuple {Source Address, Source Port, Destination Address, Destination Port}. Thus, in order to attack a TCP connection, an attacker needs to know or be able to guess each of the values in that four-tuple. Assuming the two peers of the target TCP connection are known, the Source Address and the Destination Address can be assumed to be known.

We note that in order to be able to successfully perform this attack, the attacker needs to be able to send packets with a forged source address. This means that the attacker cannot be located in a network where techniques like ingress filtering

[RFC2827] or source address validation [RFC7039] are deployed. However, ingress filtering is not as widely implemented as one would expect, and hence cannot be relied upon as a mitigation for this kind of attack.

Assuming the attacker knows the application protocol for which the TCP connection is being employed, the server-side port can also be assumed to be known. Finally, the client-side port will generally not be known, and will need to be guessed by the attacker. The chances of an attacker guessing the client-side port will depend on the ephemeral port range employed by the client, and whether the client implements port randomization [RFC6056].

Assuming TCP sequence number randomization is in place (see e.g. [RFC6528]), an attacker would have to blindly guess a valid TCP sequence number. That is,

$$\text{RCV.NXT} \leq \text{SEG.SEQ} < \text{RCV.NXT} + \text{RCV.WND} \text{ or } \text{RCV.NXT} \leq \text{SEG.SEQ} + \text{SEG.LEN} - 1 < \text{RCV.NXT} + \text{RCV.WND}$$

As a result, the chances of an attacker to succeed will depend on the TCP receive window size at the target TCP peer.

We note that automatic TCP buffer tuning mechanisms have been become common for popular TCP implementations, and hence very large TCP window sizes of values up to 2 MB could end up being employed by such TCP implementations.

According to [RFC0793], the Acknowledgement Number is considered valid as long as it does not acknowledge the receipt of data that has not yet been sent. That is, the following expression must be true:

$$\text{SEG.ACK} \leq \text{SND.NXT}$$

However, for implementations that support [RFC5961], the following (stricter) validation check is enforced:

$$\text{SND.UNA} - \text{SND.MAX.WND} \leq \text{SEG.ACK} \leq \text{SND.NXT}$$

Finally, in order for the address identifier to be valid, the only requirement is that it needs to be different than the ones already being used by Host A in that MPTCP session, so a random identifier is likely to work.

Given that a large number of factors affect the chances of an attacker of successfully performing the aforementioned off-path attacks, we provide two general expressions for the expected number of packets the attacker needs to send to succeed in the attack: one

for MPTCP implementations that support [RFC5961], and another for MPTCP implementations that do not.

Implementations that do not support RFC 5961

$$\text{Packets} = (2^{32}/(\text{RCV_WND})) * 2 * \text{EPH_PORT_SIZE}/2 * 1/\text{MSS}$$

Where the new :

Packets:

Maximum number of packets required to successfully perform an off-path (blind) attack.

RCV_WND:

TCP receive window size (RCV.WND) at the target node.

EPH_PORT_SIZE:

Number of ports comprising the ephemeral port range at the "client" system.

MSS:

Maximum Segment Size, assuming the attacker will send full segments to maximize the chances to get a hit.

Notes:

The value "2³²" represents the size of the TCP sequence number space.

The value "2" accounts for 2 different ACK numbers (separated by 2³¹) that should be employed to make sure the ACK number is valid.

The following table contains some sample results for the number of required packets, based on different values of RCV_WND and EPH_PORT_SIZE for a MSS of 1500 bytes.

Ports \ Win	16 KB	128 KB	256 KB	2048 KB
4000	699050	87381	43690	5461
10000	1747626	218453	109226	13653
50000	8738133	1092266	546133	68266

Table 1: Max. Number of Packets for Successful Attack

Implementations that do support RFC 5961

$$\text{Packets} = (2^{32}/(\text{RCV_WND})) * (2^{32}/(2 * \text{SND_MAX_WND})) * \text{EPH_PORT_SIZE}/2 * 1/\text{MSS}$$

Where:

Packets:

Maximum number of packets required to successfully perform an off-path (blind) attack.

RCV_WND:

TCP receive window size (RCV.WND) at the target MPTCP endpoint.

SND_MAX_WND:

Maximum TCP send window size ever employed by the target MPTCP end-point (SND.MAX.WND).

EPH_PORT_SIZE:

Number of ports comprising the ephemeral port range at the "client" system.

Notes:

The value "2³²" represents the size of the TCP sequence number space.

The parameter "SND_MAX_WND" is specified in [RFC5961].

The value "2*SND_MAX_WND" results from the expression "SND.NXT - SND.UNA - MAX.SND.WND", assuming that, for connections that perform bulk data transfers, "SND.NXT - SND.UNA == MAX.SND.WND". If an attacker targets a TCP endpoint that is not actively transferring data, "2 * SND_MAX_WND" would become "SND_MAX_WND" (and hence a successful attack would typically require more packets).

The following table contains some sample results for the number of required packets, based on different values of RCV_WND, SND_MAX_WND, and EPH_PORT_SIZE. For these implementations, only a limited number of sample results are provided, just as an indication of how [RFC5961] increases the difficulty of performing these attacks.

Ports \ Win	16 KB	128 KB	256 KB	2048 KB
4000	45812984490	715827882	178956970	2796202

Table 2: Max. Number of Packets for Successful Attack

Note:

In the aforementioned table, all values are computed with RCV_WND equal to SND_MAX_WND.

2.1. Possible security enhancements to prevent this attack

1. To include the token of the connection in the ADD_ADDR option. This would make it harder for the attacker to launch the attack, since he needs to either eavesdrop the token (so this can no longer be a blind attack) or to guess it, but a random 32 bit number is not so easy to guess. However, this would imply that any eavesdropper that is able to see the token, would be able to launch this attack. This solution then increases the vulnerability window against eavesdroppers from the initial 3-way handshake for the MPTCP session to any exchange of the ADD_ADDR messages.
2. To include the HMAC of the address contained in the ADD_ADDR option. The key used for the HMAC is the concatenation of the key of the receiver and the key of the sender (in the same way they are used for generating the HMAC of the MP_JOIN message). This makes it much more secure, since it requires the attacker to have both keys (either by eavesdropping it in the first exchange or by guessing it). Because this solution relies on the key used in the MPTCP session, the protection of this solution would increase if new key generation methods are defined for MPTCP (e.g. using SSL keys as has been proposed).
3. To include the destination address of the SYN packet in the HMAC of the MP_JOIN message. As the attacker requires to change the destination address to perform the described attack, protecting it would prevent the attack. It wouldn't allow hosts behind NATs to be reached by an address in the ADD_ADDR option, even with static NAT bindings (like a web server at home).

Out of the options described above, option 2 is recommended as it achieves a higher security level while preserving the required functionality (i.e. NAT compatibility).

3. DoS attack on MP_JOIN

Summary of the attack:

Type of attack: MPTCP Denial-of-Service attack, preventing the hosts from creating new subflows.

Type of attacker: Off-path, active attacker

Description:

As currently specified, the initial SYN+MP_JOIN message of the 3-way handshake for additional subflows creates state in the host receiving the message. This is because the SYN+MP_JOIN contains the 32-bit token that allows the receiver to identify the MPTCP-session and the 32-bit random nonce, used in the HMAC calculation. As this information is not resent in the third ACK of the 3-way handshake, a host must create state upon reception of a SYN+MP_JOIN.

Assume that there exists an MPTCP-session between host A and host B, with token Ta and Tb. An attacker, sending a SYN+MP_JOIN to host B, with the valid token Tb, will trigger the creation of state on host B. The number of these half-open connections a host can store per MPTCP-session is limited by a certain number, and it is implementation-dependent. The attacker can simply exhaust this limit by sending multiple SYN+MP_JOINS with different 5-tuples. The (possibly forged) source address of the attack packets will typically correspond to an address that is not in use, or else the SYN/ACK sent by Host B would elicit a RST from the impersonated node, thus removing the corresponding state at Host B. Further discussion of traditional SYN-flood attacks and common mitigations can be found in [RFC4987]

This effectively prevents the host A from sending any more SYN+MP_JOINS to host B, as the number of acceptable half-open connections per MPTCP-session on host B has been exhausted.

The attacker needs to know the token Tb in order to perform the described attack. This can be achieved if it is a Partial-time On-path eavesdropper, observing the 3-way handshake of the establishment of an additional subflow between host A and host B. If the attacker is never on-path, it has to guess the 32-bit token.

3.1. Possible security enhancements to prevent this attack

The third packet of the 3-way handshake could be extended to contain also the 32-bit token and the random nonce that has been sent in the SYN+MP_JOIN. Further, host B will have to generate its own random nonce in a reproducible fashion (e.g., a Hash of the 5-tuple + initial sequence-number + local secret). This will allow host B to reply to a SYN+MP_JOIN without having to create state. Upon the reception of the third ACK, host B can then verify the correctness of the HMAC and create the state.

4. SYN flooding amplification

Summary of the attack:

Type of attack: The attacker can use the SYN+MP_JOIN messages to amplify the SYN flooding attack.

Type of attacker: Off-path, active attacker

Description:

SYN flooding attacks [RFC4987] use SYN messages to exhaust the server's resources and prevent new TCP connections. A common mitigation is the use of SYN cookies [RFC4987] that allow the stateless processing of the initial SYN message.

With MPTCP, the initial SYN can be processed in a stateless fashion using the aforementioned SYN cookies. However, as we described in the previous section, as currently specified, the SYN+MP_JOIN messages are not processed in a stateless manner. This opens a new attack vector. The attacker can now open a MPTCP session by sending a regular SYN and creating the associated state but then send as many SYN+MP_JOIN messages as supported by the server with different source address source port combinations, consuming server's resources without having to create state in the attacker. This is an amplification attack, where the cost on the attacker side is only the cost of the state associated with the initial SYN while the cost on the server side is the state for the initial SYN plus all the state associated to all the following SYN+MP_JOIN.

4.1. Possible security enhancements to prevent this attack

1. The solution described for the previous DoS attack on MP_JOIN would also prevent this attack.
2. Limiting the number of half open subflows to a low number (e.g. 3 subflows) would also limit the impact of this attack.

5. Eavesdropper in the initial handshake

Summary of the attack

Type of attack: An eavesdropper present in the initial handshake where the keys are exchanged can hijack the MPTCP session at any time in the future.

Type of attacker: a Partial-time On-path eavesdropper

Description:

In this case, the attacker is present along the path when the initial 3-way handshake takes place, and therefore is able to learn the keys

used in the MPTCP session. This allows the attacker to move away from the MPTCP session path and still be able to hijack the MPTCP session in the future. This vulnerability was readily identified at the moment of the design of the MPTCP security solution and the threat was considered acceptable.

5.1. Possible security enhancements to prevent this attack

There are many techniques that can be used to prevent this attack and each of them represents different tradeoffs. At this point, we limit ourselves to enumerate them and provide useful pointers.

1. Use of hash-chains. The use of hash chains for MPTCP has been explored in [hash-chains]
2. Use of SSL keys for MPTCP security as described in [I-D.paasch-mptcp-ssl]
3. Use of Cryptographically-Generated Addresses (CGAs) for MPTCP security. CGAs [RFC3972] have been used in the past to secure multi addressed protocols like SHIM6 [RFC5533].
4. Use of TCPCrypt [I-D.bittau-tcp-crypt]
5. Use DNSSEC. DNSSEC has been proposed to secure the Mobile IP protocol [dnssec]

6. SYN/JOIN attack

Summary of the attack

Type of attack: An attacker that can intercept the SYN/JOIN message can alter the source address being added.

Type of attacker: a Partial-time On-path eavesdropper

Description:

The attacker is present along the path when the SYN/JOIN exchange takes place, and this allows the attacker to add any new address it wants to by simply substituting the source address of the SYN/JOIN packet for one it chooses. This vulnerability was readily identified at the moment of the design of the MPTCP security solution and the threat was considered acceptable.

6.1. Possible security enhancements to prevent this attack

It should be noted that this vulnerability is fundamental due to the NAT support requirement. In other words, MPTCP must work through NATs in order to be deployable in the current Internet. NAT behavior is unfortunately indistinguishable from this attack. It is impossible to secure the source address, since doing so would prevent MPTCP to work through NATs. This basically implies that the solution cannot rely on securing the address. A more promising approach would be then to look into securing the payload exchanged, limiting the impact that the attack would have in the communication (e.g. TCPCrypt [I-D.bittau-tcp-crypt] or similar).

7. Recommendations

Current MPTCP specification [RFC6824] is experimental. There is an ongoing effort to move it to Standards track. We believe that the work on MPTCP security should follow two threads:

- o The work on improving MPTCP security so that is enough to become a Standard Track document.
- o The work on analyzing possible additional security enhancements to provide a more secure version of MPTCP.

We will expand on these two next.

MPTCP security for a Standard Track specification.

We believe that in order for MPTCP to progress to Standard Track, the ADD_ADDR attack must be addressed. We believe that the solution that should be adopted in order to deal with this attack is to include an HMAC to the ADD_ADDR message (with the address being added used as input to the HMAC, as well as the key). This would make the ADD_ADDR message as secure as the JOIN message. In addition, this implies that if we implement a more secure way to create the key used in the MPTCP connection, then the security of both the MP_JOIN and the ADD_ADDR messages is automatically improved (since both use the same key in the HMAC).

We believe that this is enough for MPTCP to progress as a Standard track document, because the security level is similar to single path TCP, as results from our previous analysis. Moreover, the security level achieved with these changes is exactly the same as other Standard Track documents. In particular, this would be the same security level as SCTP with dynamic addresses as defined in [RFC5061]. The Security Consideration section of RFC5061 (which is a Standard Track document) reads:

The addition and or deletion of an IP address to an existing association does provide an additional mechanism by which existing associations can be hijacked. Therefore, this document requires the use of the authentication mechanism defined in [RFC4895] to limit the ability of an attacker to hijack an association. Hijacking an association by using the addition and deletion of an IP address is only possible for an attacker who is able to intercept the initial two packets of the association setup when the SCTP-AUTH extension is used without pre-shared keys. If such a threat is considered a possibility, then the [RFC4895] extension must be used with a preconfigured shared endpoint pair key to mitigate this threat.

This is the same security level that would be achieved by MPTCP plus the ADD_ADDR security measure recommended.

7.1. Security enhancements for MPTCP

We also believe that is worthwhile exploring alternatives to secure MPTCP. As we identified earlier, the problem is securing JOIN messages is fundamentally incompatible with NAT support, so it is likely that a solution to this problem involves the protection of the data itself. Exploring the integration of MPTCP and approaches like TCPCrypt [I-D.bittau-tcp-crypt] or integration with SSL seem promising venues.

8. Security considerations

This whole document is about security considerations for MPTCP.

9. IANA Considerations

There are no IANA considerations in this memo.

10. Acknowledgments

We would like to thank Mark Handley for his comments on the attacks and countermeasures discussed in this document. thanks to Alissa Copper, Phil Eardley, Yoshifumi Nishida, Barry Leiba, Stephen Farrell, Stefan Winter for the comments and review. Marcelo Bagnulo, Christophe Paasch, Oliver Bonaventure and Costin Raiciu are partially funded by the EU Trilogy 2 project.

11. References

11.1. Normative References

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005.
- [RFC5961] Ramaiah, A., Stewart, R., and M. Dalal, "Improving TCP's Robustness to Blind In-Window Attacks", RFC 5961, August 2010.
- [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-Protocol Port Randomization", BCP 156, RFC 6056, January 2011.
- [RFC6528] Gont, F. and S. Bellovin, "Defending against Sequence Number Attacks", RFC 6528, February 2012.
- [RFC5061] Stewart, R., Xie, Q., Tuexen, M., Maruyama, S., and M. Kozuka, "Stream Control Transmission Protocol (SCTP) Dynamic Address Reconfiguration", RFC 5061, September 2007.
- [RFC4895] Tuexen, M., Stewart, R., Lei, P., and E. Rescorla, "Authenticated Chunks for the Stream Control Transmission Protocol (SCTP)", RFC 4895, August 2007.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013.

11.2. Informative References

- [RFC6181] Bagnulo, M., "Threat Analysis for TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6181, March 2011.
- [RFC5533] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, June 2009.
- [RFC4960] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.
- [RFC6275] Perkins, C., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, July 2011.

- [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [RFC7039] Wu, J., Bi, J., Bagnulo, M., Baker, F., and C. Vogt, "Source Address Validation Improvement (SAVI) Framework", RFC 7039, October 2013.
- [RFC4987] Eddy, W., "TCP SYN Flooding Attacks and Common Mitigations", RFC 4987, August 2007.
- [I-D.paasch-mptcp-ssl]
Paasch, C. and O. Bonaventure, "Securing the MultiPath TCP handshake with external keys", draft-paasch-mptcp-ssl-00 (work in progress), October 2012.
- [I-D.bittau-tcp-crypt]
Bittau, A., Boneh, D., Hamburg, M., Handley, M., Mazieres, D., and Q. Slack, "Cryptographic protection of TCP Streams (tcpcrypt)", draft-bittau-tcp-crypt-04 (work in progress), February 2014.
- [hash-chains]
Diez, J., Bagnulo, M., Valera, F., and I. Vidal, "Security for multipath TCP: a constructive approach", International Journal of Internet Protocol Technology 6, 2011.
- [dnssec] Kukec, A., Bagnulo, M., Ayaz, S., Bauer, C., and W. Eddy, "OAM-DNSSEC: Route Optimization for Aeronautical Mobility using DNSSEC", 4th ACM International Workshop on Mobility in the Evolving Internet Architecture MobiArch 2009, 2009.

Authors' Addresses

Marcelo Bagnulo
Universidad Carlos III de Madrid
Av. Universidad 30
Leganes, Madrid 28911
SPAIN

Phone: 34 91 6249500
Email: marcelo@it.uc3m.es
URI: <http://www.it.uc3m.es>

Christoph Paasch
UCLouvain
Place Sainte Barbe, 2
Louvain-la-Neuve, 1348
Belgium

Email: christoph.paasch@uclouvain.be

Fernando Gont
SI6 Networks / UTN-FRH
Evaristo Carriego 2644
Haedo, Provincia de Buenos Aires 1706
Argentina

Phone: +54 11 4650 8472
Email: fgont@si6networks.com
URI: <http://www.si6networks.com>

Olivier Bonaventure
UCLouvain
Place Sainte Barbe, 2
Louvain-la-Neuve, 1348
Belgium

Email: olivier.bonaventure@uclouvain.be

Costin Raiciu
Universitatea Politehnica Bucuresti
Splaiul Independentei 313a
Bucuresti
Romania

Email: costin.raiciu@cs.pub.ro

MPTCP Working Group
Internet-Draft
Intended status: Informational
Expires: April 30, 2017

O. Bonaventure
UCLouvain
C. Paasch
Apple, Inc.
G. Detal
Tessares
October 27, 2016

Use Cases and Operational Experience with Multipath TCP
draft-ietf-mptcp-experience-07

Abstract

This document discusses both use cases and operational experience with Multipath TCP in real networks. It lists several prominent use cases where Multipath TCP has been considered and is being used. It also gives insight to some heuristics and decisions that have helped to realize these use cases and suggests possible improvements.

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 April 30, 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 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. Use cases	5
2.1. Datacenters	5
2.2. Cellular/WiFi Offload	5
2.3. Multipath TCP proxies	9
3. Operational Experience	11
3.1. Middlebox interference	11
3.2. Congestion control	13
3.3. Subflow management	13
3.4. Implemented subflow managers	14
3.5. Subflow destination port	16
3.6. Closing subflows	17
3.7. Packet schedulers	18
3.8. Segment size selection	19
3.9. Interactions with the Domain Name System	19
3.10. Captive portals	20
3.11. Stateless webservers	21
3.12. Loadbalanced server farms	22
4. IANA Considerations	23
5. Security Considerations	24
6. Acknowledgements	26
7. References	27
7.1. Normative References	27
7.2. Informative References	27
Appendix A. Changelog	34
Authors' Addresses	36

1. Introduction

Multipath TCP was specified in [RFC6824] and five independent implementations have been developed. As of November 2016, Multipath TCP has been or is being implemented on the following platforms:

- o Linux kernel [MultipathTCP-Linux]
- o Apple iOS and macOS [Apple-MPTCP]
- o Citrix load balancers
- o FreeBSD [FreeBSD-MPTCP]
- o Oracle Solaris

The first three implementations are known to interoperate. Three of these implementations are open-source (Linux kernel, FreeBSD and Apple's iOS and macOS). Apple's implementation is widely deployed.

Since the publication of [RFC6824] as an experimental RFC, experience has been gathered by various network researchers and users about the operational issues that arise when Multipath TCP is used in today's Internet.

When the MPTCP working group was created, several use cases for Multipath TCP were identified [RFC6182]. Since then, other use cases have been proposed and some have been tested and even deployed. We describe these use cases in Section 2.

Section 3 focuses on the operational experience with Multipath TCP. Most of this experience comes from the utilization of the Multipath TCP implementation in the Linux kernel [MultipathTCP-Linux]. This open-source implementation has been downloaded and is used by thousands of users all over the world. Many of these users have provided direct or indirect feedback by writing documents (scientific articles or blog messages) or posting to the mptcp-dev mailing list (see <https://listes-2.sipr.ucl.ac.be/sympa/arc/mptcp-dev>). This Multipath TCP implementation is actively maintained and continuously improved. It is used on various types of hosts, ranging from smartphones or embedded routers to high-end servers.

The Multipath TCP implementation in the Linux kernel is not, by far, the most widespread deployment of Multipath TCP. Since September 2013, Multipath TCP is also supported on smartphones and tablets since iOS7 [IOS7]. There are likely hundreds of millions of Multipath TCP enabled devices. This Multipath TCP implementation is currently only used to support the Siri voice recognition/control

application. Some lessons learned from this deployment are described in [IETFJ].

Section 3 is organized as follows. Supporting the middleboxes was one of the difficult issues in designing the Multipath TCP protocol. We explain in Section 3.1 which types of middleboxes the Linux Kernel implementation of Multipath TCP supports and how it reacts upon encountering these. Section 3.2 summarizes the MPTCP specific congestion controls that have been implemented. Section 3.3 to Section 3.7 discuss heuristics and issues with respect to subflow management as well as the scheduling across the subflows. Section 3.8 explains some problems that occurred with subflows having different Maximum Segment Size (MSS) values. Section 3.9 presents issues with respect to content delivery networks and suggests a solution to this issue. Finally, Section 3.10 documents an issue with captive portals where MPTCP will behave sub optimally.

2. Use cases

Multipath TCP has been tested in several use cases. There is already an abundant scientific literature on Multipath TCP [MPTCPBIB]. Several of the papers published in the scientific literature have identified possible improvements that are worth being discussed here.

2.1. Datacenters

A first, although initially unexpected, documented use case for Multipath TCP has been in datacenters [HotNets][SIGCOMM11]. Today's datacenters are designed to provide several paths between single-homed servers. The multiplicity of these paths comes from the utilization of Equal Cost Multipath (ECMP) and other load balancing techniques inside the datacenter. Most of the deployed load balancing techniques in datacenters rely on hashes computed over the five tuple. Thus all packets from the same TCP connection follow the same path and so are not reordered. The results in [HotNets] demonstrate by simulations that Multipath TCP can achieve a better utilization of the available network by using multiple subflows for each Multipath TCP session. Although [RFC6182] assumes that at least one of the communicating hosts has several IP addresses, [HotNets] demonstrates that Multipath TCP is beneficial when both hosts are single-homed. This idea is analyzed in more details in [SIGCOMM11] where the Multipath TCP implementation in the Linux kernel is modified to be able to use several subflows from the same IP address. Measurements in a public datacenter show the quantitative benefits of Multipath TCP [SIGCOMM11] in this environment.

Although ECMP is widely used inside datacenters, this is not the only environment where there are different paths between a pair of hosts. ECMP and other load balancing techniques such as Link Aggregation Groups (LAG) are widely used in today's networks and having multiple paths between a pair of single-homed hosts is becoming the norm instead of the exception. Although these multiple paths have often the same cost (from an IGP metrics viewpoint), they do not necessarily have the same performance. For example, [IMC13c] reports the results of a long measurement study showing that load balanced Internet paths between that same pair of hosts can have huge delay differences.

2.2. Cellular/WiFi Offload

A second use case that has been explored by several network researchers is the cellular/WiFi offload use case. Smartphones or other mobile devices equipped with two wireless interfaces are a very common use case for Multipath TCP. In September 2015, this is also the largest deployment of Multipath-TCP enabled devices [IOS7]. It

has been briefly discussed during IETF88 [ietf88], but there is no published paper or report that analyses this deployment. For this reason, we only discuss published papers that have mainly used the Multipath TCP implementation in the Linux kernel for their experiments.

The performance of Multipath TCP in wireless networks was briefly evaluated in [NSDI12]. One experiment analyzes the performance of Multipath TCP on a client with two wireless interfaces. This evaluation shows that when the receive window is large, Multipath TCP can efficiently use the two available links. However, if the window becomes smaller, then packets sent on a slow path can block the transmission of packets on a faster path. In some cases, the performance of Multipath TCP over two paths can become lower than the performance of regular TCP over the best performing path. Two heuristics, reinjection and penalization, are proposed in [NSDI12] to solve this identified performance problem. These two heuristics have since been used in the Multipath TCP implementation in the Linux kernel. [CONEXT13] explored the problem in more detail and revealed some other scenarios where Multipath TCP can have difficulties in efficiently pooling the available paths. Improvements to the Multipath TCP implementation in the Linux kernel are proposed in [CONEXT13] to cope with some of these problems.

The first experimental analysis of Multipath TCP in a public wireless environment was presented in [Cellnet12]. These measurements explore the ability of Multipath TCP to use two wireless networks (real WiFi and 3G networks). Three modes of operation are compared. The first mode of operation is the simultaneous use of the two wireless networks. In this mode, Multipath TCP pools the available resources and uses both wireless interfaces. This mode provides fast handover from WiFi to cellular or the opposite when the user moves. Measurements presented in [CACM14] show that the handover from one wireless network to another is not an abrupt process. When a host moves, there are regions where the quality of one of the wireless networks is weaker than the other, but the host considers this wireless network to still be up. When a mobile host enters such regions, its ability to send packets over another wireless network is important to ensure a smooth handover. This is clearly illustrated from the packet trace discussed in [CACM14].

Many cellular networks use volume-based pricing and users often prefer to use unmetered WiFi networks when available instead of metered cellular networks. [Cellnet12] implements support for the MP_PRIO option to explore two other modes of operation.

In the backup mode, Multipath TCP opens a TCP subflow over each interface, but the cellular interface is configured in backup mode.

This implies that data flows only over the WiFi interface when both interfaces are considered to be active. If the WiFi interface fails, then the traffic switches quickly to the cellular interface, ensuring a smooth handover from the user's viewpoint [Cellnet12]. The cost of this approach is that the WiFi and cellular interfaces are likely to remain active all the time since all subflows are established over the two interfaces.

The single-path mode is slightly different. This mode benefits from the break-before-make capability of Multipath TCP. When an MPTCP session is established, a subflow is created over the WiFi interface. No packet is sent over the cellular interface as long as the WiFi interface remains up [Cellnet12]. This implies that the cellular interface can remain idle and battery capacity is preserved. When the WiFi interface fails, a new subflow is established over the cellular interface in order to preserve the established Multipath TCP sessions. Compared to the backup mode described earlier, measurements reported in [Cellnet12] indicate that this mode of operation is characterized by a throughput drop while the cellular interface is brought up and the subflows are reestablished.

From a protocol viewpoint, [Cellnet12] discusses the problem posed by the unreliability of the REMOVE_ADDR option and proposes a small protocol extension to allow hosts to reliably exchange this option. It would be useful to analyze packet traces to understand whether the unreliability of the REMOVE_ADDR option poses an operational problem in real deployments.

Another study of the performance of Multipath TCP in wireless networks was reported in [IMC13b]. This study uses laptops connected to various cellular ISPs and WiFi hotspots. It compares various file transfer scenarios. [IMC13b] observes that 4-path MPTCP outperforms 2-path MPTCP, especially for larger files. However, for three congestion control algorithms (LIA, OLIA and Reno - see Section 3.2), there is no significant performance difference for file sizes smaller than 4MB.

A different study of the performance of Multipath TCP with two wireless networks is presented in [INFOCOM14]. In this study the two networks had different qualities : a good network and a lossy network. When using two paths with different packet loss ratios, the Multipath TCP congestion control scheme moves traffic away from the lossy link that is considered to be congested. However, [INFOCOM14] documents an interesting scenario that is summarized hereafter.

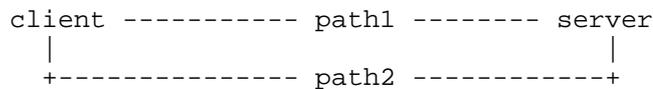


Figure 1: Simple network topology

Initially, the two paths in Figure 1 have the same quality and Multipath TCP distributes the load over both of them. During the transfer, the path2 becomes lossy, e.g. because the client moves. Multipath TCP detects the packet losses and they are retransmitted over path1. This enables the data transfer to continue over this path. However, the subflow over path2 is still up and transmits one packet from time to time. Although the N packets have been acknowledged over the first subflow (at the MPTCP level), they have not been acknowledged at the TCP level over the second subflow. To preserve the continuity of the sequence numbers over the second subflow, TCP will continue to retransmit these segments until either they are acknowledged or the maximum number of retransmissions is reached. This behavior is clearly inefficient and may lead to blocking since the second subflow will consume window space to be able to retransmit these packets. [INFOCOM14] proposes a new Multipath TCP option to solve this problem. In practice, a new TCP option is probably not required. When the client detects that the data transmitted over the second subflow has been acknowledged over the first subflow, it could decide to terminate the second subflow by sending a RST segment. If the interface associated to this subflow is still up, a new subflow could be immediately reestablished. It would then be immediately usable to send new data and would not be forced to first retransmit the previously transmitted data. As of this writing, this dynamic management of the subflows is not yet implemented in the Multipath TCP implementation in the Linux kernel.

Some studies have started to analyze the performance of Multipath TCP on smartphones with real applications. In contrast with the bulk transfers that are used by many publications, many deployed applications do not exchange huge amounts of data and mainly use small connections. [COMMAG2016] proposes a software testing framework that allows to automate Android applications to study their interactions with Multipath TCP. [PAM2016] analyses a one-month packet trace of all the packets exchanged by a dozen of smartphones used by regular users. This analysis reveals that short connections are important on smartphones and that the main benefit of using Multipath TCP on smartphones is the ability to perform seamless handovers between different wireless networks. Long connections benefit from these handovers.

2.3. Multipath TCP proxies

As Multipath TCP is not yet widely deployed on both clients and servers, several deployments have used various forms of proxies. Two families of solutions are currently being used or tested.

A first use case is when a Multipath TCP enabled client wants to use several interfaces to reach a regular TCP server. A typical use case is a smartphone that needs to use both its WiFi and its cellular interface to transfer data. Several types of proxies are possible for this use case. An HTTP proxy deployed on a Multipath TCP capable server would enable the smartphone to use Multipath TCP to access regular web servers. Obviously, this solution only works for applications that rely on HTTP. Another possibility is to use a proxy that can convert any Multipath TCP connection into a regular TCP connection. Multipath TCP-specific proxies have been proposed [HotMiddlebox13b] [HAMPEL].

Another possibility leverages the SOCKS protocol [RFC1928]. SOCKS is often used in enterprise networks to allow clients to reach external servers. For this, the client opens a TCP connection to the SOCKS server that relays it to the final destination. If both the client and the SOCKS server use Multipath TCP, but not the final destination, then Multipath TCP can still be used on the path between the clients and the SOCKS server. At IETF'93, Korea Telecom announced that they have deployed in June 2015 a commercial service that uses Multipath TCP on smartphones. These smartphones access regular TCP servers through a SOCKS proxy. This enables them to achieve throughputs of up to 850 Mbps [KT].

Measurements performed with Android smartphones [Mobicom15] show that popular applications work correctly through a SOCKS proxy and Multipath TCP enabled smartphones. Thanks to Multipath TCP, long-lived connections can be spread over the two available interfaces. However, for short-lived connections, most of the data is sent over the initial subflow that is created over the interface corresponding to the default route and the second subflow is almost not used [PAM2016].

A second use case is when Multipath TCP is used by middleboxes, typically inside access networks. Various network operators are discussing and evaluating solutions for hybrid access networks [TR-348]. Such networks arise when a network operator controls two different access network technologies, e.g. wired and cellular, and wants to combine them to improve the bandwidth offered to the endusers [I-D.lhwxz-hybrid-access-network-architecture]. Several solutions are currently investigated for such networks [TR-348]. Figure 2 shows the organization of such a network. When a client

creates a normal TCP connection, it is intercepted by the Hybrid CPE (HCPE) that converts it in a Multipath TCP connection so that it can use the available access networks (DSL and LTE in the example). The Hybrid Access Gateway (HAG) does the opposite to ensure that the regular server sees a normal TCP connection. Some of the solutions currently discussed for hybrid networks use Multipath TCP on the HCPE and the HAG. Other solutions rely on tunnels between the HCPE and the HAG [I-D.lhwxz-gre-notifications-hybrid-access].

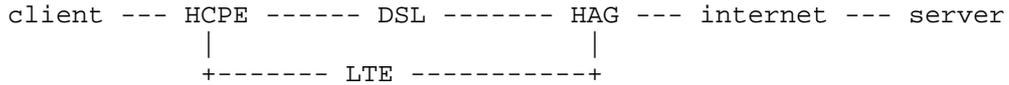


Figure 2: Hybrid Access Network

3. Operational Experience

3.1. Middlebox interference

The interference caused by various types of middleboxes has been an important concern during the design of the Multipath TCP protocol. Three studies on the interactions between Multipath TCP and middleboxes are worth discussing.

The first analysis appears in [IMC11]. This paper was the main motivation for Multipath TCP incorporating various techniques to cope with middlebox interference. More specifically, Multipath TCP has been designed to cope with middleboxes that :

- o change source or destination addresses
- o change source or destination port numbers
- o change TCP sequence numbers
- o split or coalesce segments
- o remove TCP options
- o modify the payload of TCP segments

These middlebox interferences have all been included in the MBtest suite [MBTest]. This test suite is used in [HotMiddlebox13] to verify the reaction of the Multipath TCP implementation in the Linux kernel [MultipathTCP-Linux] when faced with middlebox interference. The test environment used for this evaluation is a dual-homed client connected to a single-homed server. The middlebox behavior can be activated on any of the paths. The main results of this analysis are :

- o the Multipath TCP implementation in the Linux kernel, is not affected by a middlebox that performs NAT or modifies TCP sequence numbers
- o when a middlebox removes the MP_CAPABLE option from the initial SYN segment, the Multipath TCP implementation in the Linux kernel falls back correctly to regular TCP
- o when a middlebox removes the DSS option from all data segments, the Multipath TCP implementation in the Linux kernel falls back correctly to regular TCP

- o when a middlebox performs segment coalescing, the Multipath TCP implementation in the Linux kernel is still able to accurately extract the data corresponding to the indicated mapping
- o when a middlebox performs segment splitting, the Multipath TCP implementation in the Linux kernel correctly reassembles the data corresponding to the indicated mapping. [HotMiddlebox13] shows on figure 4 in section 3.3 a corner case with segment splitting that may lead to a desynchronization between the two hosts.

The interactions between Multipath TCP and real deployed middleboxes is also analyzed in [HotMiddlebox13] and a particular scenario with the FTP application level gateway running on a NAT is described.

Middlebox interference can also be detected by analyzing packet traces on Multipath TCP enabled servers. A closer look at the packets received on the multipath-tcp.org server [TMA2015] shows that among the 184,000 Multipath TCP connections, only 125 of them were falling back to regular TCP. These connections originated from 28 different client IP addresses. These include 91 HTTP connections and 34 FTP connections. The FTP interference is expected since Application Level Gateways used for FTP modify the TCP payload and the DSS Checksum detects these modifications. The HTTP interference appeared only on the direction from server to client and could have been caused by transparent proxies deployed in cellular or enterprise networks. A longer trace is discussed in [COMCOM2016] and similar conclusions about the middlebox interference are provided.

From an operational viewpoint, knowing that Multipath TCP can cope with various types of middlebox interference is important. However, there are situations where the network operators need to gather information about where a particular middlebox interference occurs. The tracebox software [tracebox] described in [IMC13a] is an extension of the popular traceroute software that enables network operators to check at which hop a particular field of the TCP header (including options) is modified. It has been used by several network operators to debug various middlebox interference problems. Experience with tracebox indicates that supporting the ICMP extension defined in [RFC1812] makes it easier to debug middlebox problems in IPv4 networks.

Users of the Multipath TCP implementation have reported some experience with middlebox interference. The strangest scenario has been a middlebox that accepts the Multipath TCP options in the SYN segment but later replaces Multipath TCP options with a TCP EOL option [StrangeMbox]. This causes Multipath TCP to perform a fallback to regular TCP without any impact on the application.

3.2. Congestion control

Congestion control has been an important challenge for Multipath TCP. The congestion control scheme specified for Multipath TCP is defined in [RFC6356]. A detailed description of this algorithm is provided in [NSDI11]. This congestion control scheme has been implemented in the Linux implementation of Multipath TCP. Linux uses a modular architecture to support various congestion control schemes. This architecture is applicable for both regular TCP and Multipath TCP. While the coupled congestion control scheme defined in [RFC6356] is the default congestion control scheme in the Linux implementation, other congestion control schemes have been added. The second congestion control scheme is OLIA [CONEXT12]. This congestion control scheme is also an adaptation of the NewReno single path congestion control scheme to support multiple paths. Simulations and measurements have shown that it provides some performance benefits compared to the default congestion control scheme [CONEXT12]. Measurements over a wide range of parameters reported in [CONEXT13] also indicate some benefits with the OLIA congestion control scheme. Recently, a delay-based congestion control scheme has been ported to the Multipath TCP implementation in the Linux kernel. This congestion control scheme has been evaluated by using simulations in [ICNP12] and measurements in [PaaschPhD]. The fourth congestion control scheme that has been included in the Linux implementation of Multipath TCP is the BALIA scheme that provides a better balance between TCP friendliness, responsiveness, and window oscillation [BALIA].

These different congestion control schemes have been compared in several articles. [CONEXT13] and [PaaschPhD] compare these algorithms in an emulated environment. The evaluation showed that the delay-based congestion control scheme is less able to efficiently use the available links than the three other schemes.

3.3. Subflow management

The multipath capability of Multipath TCP comes from the utilization of one subflow per path. The Multipath TCP architecture [RFC6182] and the protocol specification [RFC6824] define the basic usage of the subflows and the protocol mechanisms that are required to create and terminate them. However, there are no guidelines on how subflows are used during the lifetime of a Multipath TCP session. Most of the published experiments with Multipath TCP have been performed in controlled environments. Still, based on the experience running them and discussions on the mptcp-dev mailing list, interesting lessons have been learned about the management of these subflows.

From a subflow viewpoint, the Multipath TCP protocol is completely

symmetrical. Both the clients and the server have the capability to create subflows. However in practice the existing Multipath TCP implementations have opted for a strategy where only the client creates new subflows. The main motivation for this strategy is that often the client resides behind a NAT or a firewall, preventing passive subflow openings on the client. Although there are environments such as datacenters where this problem does not occur, as of this writing, no precise requirement has emerged for allowing the server to create new subflows.

3.4. Implemented subflow managers

The Multipath TCP implementation in the Linux kernel includes several strategies to manage the subflows that compose a Multipath TCP session. The basic subflow manager is the full-mesh. As the name implies, it creates a full-mesh of subflows between the communicating hosts.

The most frequent use case for this subflow manager is a multihomed client connected to a single-homed server. In this case, one subflow is created for each interface on the client. The current implementation of the full-mesh subflow manager is static. The subflows are created immediately after the creation of the initial subflow. If one subflow fails during the lifetime of the Multipath TCP session (e.g. due to excessive retransmissions, or the loss of the corresponding interface), it is not always reestablished. There is ongoing work to enhance the full-mesh path manager to deal with such events.

When the server is multihomed, using the full-mesh subflow manager may lead to a large number of subflows being established. For example, consider a dual-homed client connected to a server with three interfaces. In this case, even if the subflows are only created by the client, 6 subflows will be established. This may be excessive in some environments, in particular when the client and/or the server have a large number of interfaces. Implementations should limit the number of subflows that are used.

Creating subflows between multihomed clients and servers may sometimes lead to operational issues as observed by discussions on the mptcp-dev mailing list. In some cases the network operators would like to have a better control on how the subflows are created by Multipath TCP [I-D.boucadair-mptcp-max-subflow]. This might require the definition of policy rules to control the operation of the subflow manager. The two scenarios below illustrate some of these requirements.

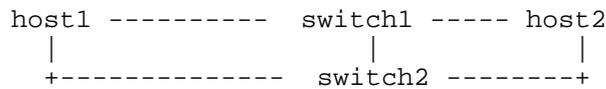


Figure 3: Simple switched network topology

Consider the simple network topology shown in Figure 3. From an operational viewpoint, a network operator could want to create two subflows between the communicating hosts. From a bandwidth utilization viewpoint, the most natural paths are host1-switch1-host2 and host1-switch2-host2. However, a Multipath TCP implementation running on these two hosts may sometimes have difficulties to obtain this result.

To understand the difficulty, let us consider different allocation strategies for the IP addresses. A first strategy is to assign two subnets : subnetA (resp. subnetB) contains the IP addresses of host1's interface to switch1 (resp. switch2) and host2's interface to switch1 (resp. switch2). In this case, a Multipath TCP subflow manager should only create one subflow per subnet. To enforce the utilization of these paths, the network operator would have to specify a policy that prefers the subflows in the same subnet over subflows between addresses in different subnets. It should be noted that the policy should probably also specify how the subflow manager should react when an interface or subflow fails.

A second strategy is to use a single subnet for all IP addresses. In this case, it becomes more difficult to specify a policy that indicates which subflows should be established.

The second subflow manager that is currently supported by the Multipath TCP implementation in the Linux kernel is the `ndiffport` subflow manager. This manager was initially created to exploit the path diversity that exists between single-homed hosts due to the utilization of flow-based load balancing techniques [SIGCOMM11]. This subflow manager creates N subflows between the same pair of IP addresses. The N subflows are created by the client and differ only in the source port selected by the client. It was not designed to be used on multihomed hosts.

A more flexible subflow manager has been proposed, implemented and evaluated in [CONEXT15]. This subflow manager exposes various kernel events to a user space daemon that decides when subflows need to be created and terminated based on various policies.

3.5. Subflow destination port

The Multipath TCP protocol relies on the token contained in the MP_JOIN option to associate a subflow to an existing Multipath TCP session. This implies that there is no restriction on the source address, destination address and source or destination ports used for the new subflow. The ability to use different source and destination addresses is key to support multihomed servers and clients. The ability to use different destination port numbers is worth discussing because it has operational implications.

For illustration, consider a dual-homed client that creates a second subflow to reach a single-homed server as illustrated in Figure 4.

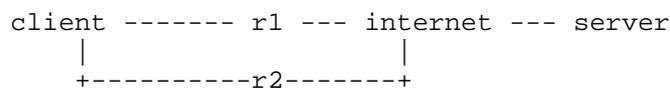


Figure 4: Multihomed-client connected to single-homed server

When the Multipath TCP implementation in the Linux kernel creates the second subflow it uses the same destination port as the initial subflow. This choice is motivated by the fact that the server might be protected by a firewall and only accept TCP connections (including subflows) on the official port number. Using the same destination port for all subflows is also useful for operators that rely on the port numbers to track application usage in their network.

There have been suggestions from Multipath TCP users to modify the implementation to allow the client to use different destination ports to reach the server. This suggestion seems mainly motivated by traffic shaping middleboxes that are used in some wireless networks. In networks where different shaping rates are associated to different destination port numbers, this could allow Multipath TCP to reach a higher performance. This behavior is valid according to the Multipath TCP specification [RFC6824]. An application could use an enhanced socket API [SOCKET] to behave in this way.

However, from an implementation point-of-view supporting different destination ports for the same Multipath TCP connection can cause some issues. A legacy implementation of a TCP stack creates a listening socket to react upon incoming SYN segments. The listening socket is handling the SYN segments that are sent on a specific port number. Demultiplexing incoming segments can thus be done solely by looking at the IP addresses and the port numbers. With Multipath TCP however, incoming SYN segments may have an MP_JOIN option with a different destination port. This means, that all incoming segments

that did not match on an existing listening-socket or an already established socket must be parsed for an eventual MP_JOIN option. This imposes an additional cost on servers, previously not existent on legacy TCP implementations.

3.6. Closing subflows

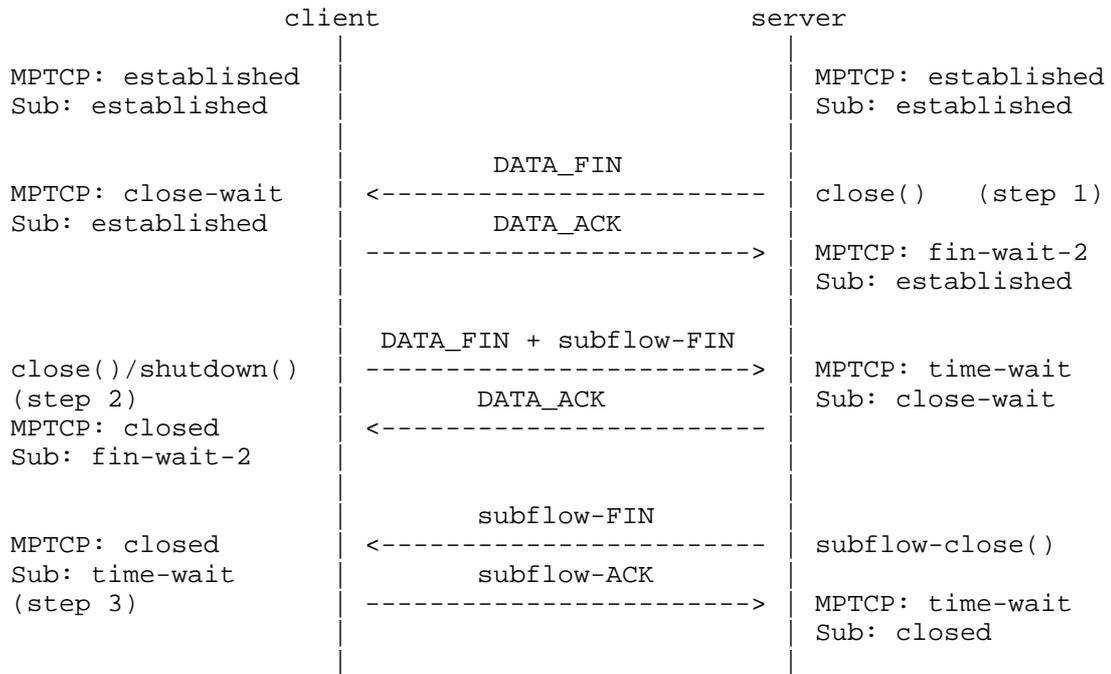


Figure 5: Multipath TCP may not be able to avoid time-wait state on the subflow (indicated as Sub in the drawing), even if enforced by the application on the client-side.

Figure 5 shows a very particular issue within Multipath TCP. Many high-performance applications try to avoid Time-Wait state by deferring the closure of the connection until the peer has sent a FIN. That way, the client on the left of Figure 5 does a passive closure of the connection, transitioning from Close-Wait to Last-ACK and finally freeing the resources after reception of the ACK of the FIN. An application running on top of a Multipath TCP enabled Linux kernel might also use this approach. The difference here is that the close() of the connection (Step 1 in Figure 5) only triggers the sending of a DATA_FIN. Nothing guarantees that the kernel is ready to combine the DATA_FIN with a subflow-FIN. The reception of the DATA_FIN will make the application trigger the closure of the

connection (step 2), trying to avoid Time-Wait state with this late closure. This time, the kernel might decide to combine the DATA_FIN with a subflow-FIN. This decision will be fatal, as the subflow's state machine will not transition from Close-Wait to Last-Ack, but rather go through Fin-Wait-2 into Time-Wait state. The Time-Wait state will consume resources on the host for at least 2 MSL (Maximum Segment Lifetime). Thus, a smart application that tries to avoid Time-Wait state by doing late closure of the connection actually ends up with one of its subflows in Time-Wait state. A high-performance Multipath TCP kernel implementation should honor the desire of the application to do passive closure of the connection and successfully avoid Time-Wait state - even on the subflows.

The solution to this problem lies in an optimistic assumption that a host doing active-closure of a Multipath TCP connection by sending a DATA_FIN will soon also send a FIN on all its subflows. Thus, the passive closer of the connection can simply wait for the peer to send exactly this FIN - enforcing passive closure even on the subflows. Of course, to avoid consuming resources indefinitely, a timer must limit the time our implementation waits for the FIN.

3.7. Packet schedulers

In a Multipath TCP implementation, the packet scheduler is the algorithm that is executed when transmitting each packet to decide on which subflow it needs to be transmitted. The packet scheduler itself does not have any impact on the interoperability of Multipath TCP implementations. However, it may clearly impact the performance of Multipath TCP sessions. The Multipath TCP implementation in the Linux kernel supports a pluggable architecture for the packet scheduler [PaaschPhD]. As of this writing, two schedulers have been implemented: round-robin and lowest-rtt-first. The second scheduler relies on the round-trip-time (rtt) measured on each TCP subflow and sends first segments over the subflow having the lowest round-trip-time. They are compared in [CSWS14]. The experiments and measurements described in [CSWS14] show that the lowest-rtt-first scheduler appears to be the best compromise from a performance viewpoint. Another study of the packet schedulers is presented in [PAMS2014]. This study relies on simulations with the Multipath TCP implementation in the Linux kernel. They compare the lowest-rtt-first with the round-robin and a random scheduler. They show some situations where the lowest-rtt-first scheduler does not perform as well as the other schedulers, but there are many scenarios where the opposite is true. [PAMS2014] notes that "it is highly likely that the optimal scheduling strategy depends on the characteristics of the paths being used."

3.8. Segment size selection

When an application performs a write/send system call, the kernel allocates a packet buffer (`sk_buff` in Linux) to store the data the application wants to send. The kernel will store at most one MSS (Maximum Segment Size) of data per buffer. As the MSS can differ amongst subflows, an MPTCP implementation must select carefully the MSS used to generate application data. The Linux kernel implementation had various ways of selecting the MSS: minimum or maximum amongst the different subflows. However, these heuristics of MSS selection can cause significant performance issues in some environment. Consider the following example. An MPTCP connection has two established subflows that respectively use a MSS of 1420 and 1428 bytes. If MPTCP selects the maximum, then the application will generate segments of 1428 bytes of data. An MPTCP implementation will have to split the segment in two (1420-byte and 8-byte) segments when pushing on the subflow with the smallest MSS. The latter segment will introduce a large overhead as for a single data segment 2 slots will be used in the congestion window (in packets) therefore reducing by roughly twice the potential throughput (in bytes/s) of this subflow. Taking the smallest MSS does not solve the issue as there might be a case where the subflow with the smallest MSS only sends a few packets therefore reducing the potential throughput of the other subflows.

The Linux implementation recently took another approach [DetailMSS]. Instead of selecting the minimum and maximum values, it now dynamically adapts the MSS based on the contribution of all the subflows to the connection's throughput. For this it computes, for each subflow, the potential throughput achieved by selecting each MSS value and by taking into account the lost space in the congestion window. It then selects the MSS that allows to achieve the highest potential throughput.

Given the prevalence of middleboxes that clamp the MSS, Multipath TCP implementations must be able to efficiently support subflows with different MSS values. The strategy described above is a possible solution to this problem.

3.9. Interactions with the Domain Name System

Multihomed clients such as smartphones can send DNS queries over any of their interfaces. When a single-homed client performs a DNS query, it receives from its local resolver the best answer for its request. If the client is multihomed, the answer in response to the DNS query may vary with the interface over which it has been sent.

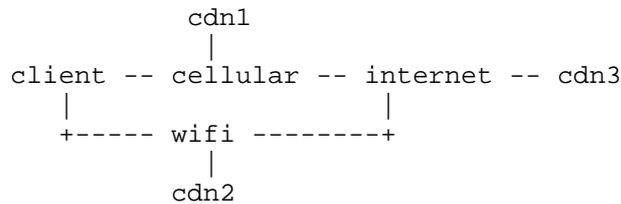


Figure 6: Simple network topology

If the client sends a DNS query over the WiFi interface, the answer will point to the cdn2 server while the same request sent over the cellular interface will point to the cdn1 server. This might cause problems for CDN providers that locate their servers inside ISP networks and have contracts that specify that the CDN server will only be accessed from within this particular ISP. Assume now that both the client and the CDN servers support Multipath TCP. In this case, a Multipath TCP session from cdn1 or cdn2 would potentially use both the cellular network and the WiFi network. Serving the client from cdn2 over the cellular interface could violate the contract between the CDN provider and the network operators. A similar problem occurs with regular TCP if the client caches DNS replies. For example the client obtains a DNS answer over the cellular interface and then stops this interface and starts to use its WiFi interface. If the client retrieves data from cdn1 over its WiFi interface, this may also violate the contract between the CDN and the network operators.

A possible solution to prevent this problem would be to modify the DNS resolution on the client. The client subnet EDNS extension defined in [RFC7871] could be used for this purpose. When the client sends a DNS query from its WiFi interface, it should also send the client subnet corresponding to the cellular interface in this request. This would indicate to the resolver that the answer should be valid for both the WiFi and the cellular interfaces (e.g., the cdn3 server).

3.10. Captive portals

Multipath TCP enables a host to use different interfaces to reach a server. In theory, this should ensure connectivity when at least one of the interfaces is active. In practice however, there are some particular scenarios with captive portals that may cause operational problems. The reference environment is shown in Figure 7.

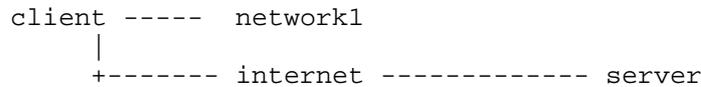


Figure 7: Issue with captive portal

The client is attached to two networks : network1 that provides limited connectivity and the entire Internet through the second network interface. In practice, this scenario corresponds to an open WiFi network with a captive portal for network1 and a cellular service for the second interface. On many smartphones, the WiFi interface is preferred over the cellular interface. If the smartphone learns a default route via both interfaces, it will typically prefer to use the WiFi interface to send its DNS request and create the first subflow. This is not optimal with Multipath TCP. A better approach would probably be to try a few attempts on the WiFi interface and then, upon failure of these attempts, try to use the second interface for the initial subflow as well.

3.11. Stateless webservers

MPTCP has been designed to interoperate with webservers that benefit from SYN-cookies to protect against SYN-flooding attacks [RFC4987]. MPTCP achieves this by echoing the keys negotiated during the MP_CAPABLE handshake in the third ACK of the 3-way handshake. Reception of this third ACK then allows the server to reconstruct the state specific to MPTCP.

However, one caveat to this mechanism is the non-reliable nature of the third ACK. Indeed, when the third ACK gets lost, the server will not be able to reconstruct the MPTCP-state. MPTCP will fallback to regular TCP in this case. This is in contrast to regular TCP. When the client starts sending data, the first data segment also includes the SYN-cookie, which allows the server to reconstruct the TCP-state. Further, this data segment will be retransmitted by the client in case it gets lost and thus is resilient against loss. MPTCP does not include the keys in this data segment and thus the server cannot reconstruct the MPTCP state.

This issue might be considered as a minor one for MPTCP. Losing the third ACK should only happen when packet loss is high. However, when packet-loss is high MPTCP provides a lot of benefits as it can move traffic away from the lossy link. It is undesirable that MPTCP has a higher chance to fall back to regular TCP in those lossy environments.

[I-D.paasch-mptcp-syncookies] discusses this issue and suggests a

modified handshake mechanism that ensures reliable delivery of the MP_CAPABLE, following the 3-way handshake. This modification will make MPTCP reliable, even in lossy environments when servers need to use SYN-cookies to protect against SYN-flooding attacks.

3.12. Loadbalanced server farms

Large-scale server farms typically deploy thousands of servers behind a single virtual IP (VIP). Steering traffic to these servers is done through layer-4 load balancers that ensure that a TCP-flow will always be routed to the same server [Presto08].

As Multipath TCP uses multiple different TCP subflows to steer the traffic across the different paths, load balancers need to ensure that all these subflows are routed to the same server. This implies that the load balancers need to track the MPTCP-related state, allowing them to parse the token in the MP_JOIN and assign those subflows to the appropriate server. However, server farms typically deploy several load balancers for reliability and capacity reasons. As a TCP subflow might get routed to any of these load balancers, they would need to synchronize the MPTCP-related state - a solution that is not feasible at large scale.

The token (carried in the MP_JOIN) contains the information indicating which MPTCP-session the subflow belongs to. As the token is a hash of the key, servers are not able to generate the token in such a way that the token can provide the necessary information to the load balancers, which would allow them to route TCP subflows to the appropriate server. [I-D.paasch-mptcp-loadbalancer] discusses this issue in detail and suggests two alternative MP_CAPABLE handshakes to overcome these.

4. IANA Considerations

There are no IANA considerations in this informational document.

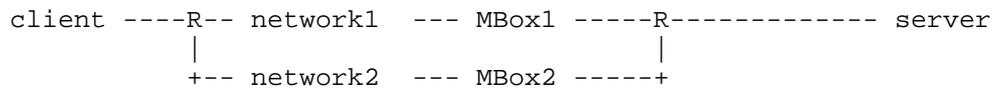


Figure 9: Interactions between load balancing and security
Middleboxes

In these two cases, it is possible for an attacker to evade some security measures operating on the TCP byte stream and implemented on the middleboxes by controlling the bytes that are actually sent over each subflow and there are tools that ease those kinds of evasion [PZ15] [PT14]. This is not a security issue for Multipath TCP itself since Multipath TCP behaves correctly. However, this demonstrates the difficulty of enforcing security policies by relying only on on-path middleboxes instead of enforcing them directly on the endpoints.

6. Acknowledgements

This work was partially supported by the FP7-Trilogy2 project. We would like to thank all the implementers and users of the Multipath TCP implementation in the Linux kernel. This document has benefited from the comments of John Ronan, Yoshifumi Nishida, Phil Eardley, Jaehyun Hwang, Mirja Kuehlewind, Benoit Claise, Jari Arkko, Qin Wu, Spencer Dawkins and Ben Campbell.

7. References

7.1. Normative References

- [RFC6182] Ford, A., Raiciu, C., Handley, M., Barre, S., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development", RFC 6182, DOI 10.17487/RFC6182, March 2011, <<http://www.rfc-editor.org/info/rfc6182>>.
- [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>>.

7.2. Informative References

- [Apple-MPTCP]
Apple, Inc, ., "iOS - Multipath TCP Support in iOS 7", n.d., <<https://support.apple.com/en-us/HT201373>>.
- [BALIA] Peng, Q., Walid, A., Hwang, J., and S. Low, "Multipath TCP Analysis, Design, and Implementation", IEEE/ACM Trans. on Networking, Vol. 24, No. 1 , 2016.
- [CACM14] Paasch, C. and O. Bonaventure, "Multipath TCP", Communications of the ACM, 57(4):51-57 , April 2014, <<http://inl.info.ucl.ac.be/publications/multipath-tcp>>.
- [COMCOM2016]
Tran, V., De Coninck, Q., Hesmans, B., Sadre, R., and O. Bonaventure, "Observing real Multipath TCP traffic", Computer Communications , April 2016, <<http://inl.info.ucl.ac.be/publications/observing-real-multipath-tcp-traffic>>.
- [COMMAG2016]
De Coninck, Q., Baerts, M., Hesmans, B., and O. Bonaventure, "Observing Real Smartphone Applications over Multipath TCP", IEEE Communications Magazine , March 2016, <<http://inl.info.ucl.ac.be/publications/observing-real-smartphone-applications-over-multipath-tcp>>.
- [CONEXT12]
Khalili, R., Gast, N., Popovic, M., Upadhyay, U., and J. Leboudec, "MPTCP is not pareto-optimal performance issues and a possible solution", Proceedings of the 8th international conference on Emerging networking

experiments and technologies (CoNEXT12) , 2012.

[CONEXT13]

Paasch, C., Khalili, R., and O. Bonaventure, "On the Benefits of Applying Experimental Design to Improve Multipath TCP", Conference on emerging Networking EXperiments and Technologies (CoNEXT) , December 2013, <<http://inl.info.ucl.ac.be/publications/benefits-applying-experimental-design-improve-multipath-tcp>>.

[CONEXT15]

Hesmans, B., Detal, G., Barre, S., Bauduin, R., and O. Bonaventure, "SMAPP - Towards Smart Multipath TCP-enabled Applications", Proc. Conext 2015, Heidelberg, Germany , December 2015, <<http://inl.info.ucl.ac.be/publications/smapp-towards-smart-multipath-tcp-enabled-applications>>.

[CSWS14]

Paasch, C., Ferlin, S., Alay, O., and O. Bonaventure, "Experimental Evaluation of Multipath TCP Schedulers", SIGCOMM CSWS2014 workshop , August 2014.

[Cellnet12]

Paasch, C., Detal, G., Duchene, F., Raiciu, C., and O. Bonaventure, "Exploring Mobile/WiFi Handover with Multipath TCP", ACM SIGCOMM workshop on Cellular Networks (Cellnet12) , 2012, <<http://inl.info.ucl.ac.be/publications/exploring-mobilewifi-handover-multipath-tcp>>.

[DetalMSS]

Detal, G., "Adaptive MSS value", Post on the mptcp-dev mailing list , September 2014, <<https://listes-2.sipr.ucl.ac.be/sympa/arc/mptcp-dev/2014-09/msg00130.html>>.

[FreeBSD-MPTCP]

Williams, N., "Multipath TCP For FreeBSD Kernel Patch v0.5", n.d., <<http://caia.swin.edu.au/urp/newtcp/mptcp>>.

[HAMPEL]

Hempel, G., Rana, A., and T. Klein, "Seamless TCP mobility using lightweight MPTCP proxy", Proceedings of the 11th ACM international symposium on Mobility management and wireless access (MobiWac '13). , 2013.

[HotMiddlebox13]

Hesmans, B., Duchene, F., Paasch, C., Detal, G., and O. Bonaventure, "Are TCP Extensions Middlebox-proof?", CoNEXT workshop HotMiddlebox , December 2013, <<http://>>

inl.info.ucl.ac.be/publications/are-tcp-extensions-middlebox-proof>.

[HotMiddlebox13b]

Detal, G., Paasch, C., and O. Bonaventure, "Multipath in the Middle(Box)", HotMiddlebox'13 , December 2013, <<http://inl.info.ucl.ac.be/publications/multipath-middlebox>>.

[HotNets]

Raiciu, C., Pluntke, C., Barre, S., Greenhalgh, A., Wischik, D., and M. Handley, "Data center networking with multipath TCP", Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks (Hotnets-IX) , 2010, <<http://doi.acm.org/10.1145/1868447.1868457>>.

[I-D.boucadair-mptcp-max-subflow]

Boucadair, M. and C. Jacquenet, "Negotiating the Maximum Number of Multipath TCP (MPTCP) Subflows", draft-boucadair-mptcp-max-subflow-02 (work in progress), May 2016.

[I-D.lhwxyz-gre-notifications-hybrid-access]

Leymann, N., Heidemann, C., Wasserman, M., Xue, L., and M. Zhang, "GRE Notifications for Hybrid Access", draft-lhwxyz-gre-notifications-hybrid-access-01 (work in progress), January 2015.

[I-D.lhwxyz-hybrid-access-network-architecture]

Leymann, N., Heidemann, C., Wasserman, M., Xue, L., and M. Zhang, "Hybrid Access Network Architecture", draft-lhwxyz-hybrid-access-network-architecture-02 (work in progress), January 2015.

[I-D.paasch-mptcp-loadbalancer]

Paasch, C., Greenway, G., and A. Ford, "Multipath TCP behind Layer-4 loadbalancers", draft-paasch-mptcp-loadbalancer-00 (work in progress), September 2015.

[I-D.paasch-mptcp-syncookies]

Paasch, C., Biswas, A., and D. Haas, "Making Multipath TCP robust for stateless webservers", draft-paasch-mptcp-syncookies-02 (work in progress), October 2015.

[ICNP12]

Cao, Y., Xu, M., and X. Fu, "Delay-based congestion control for multipath TCP", 20th IEEE International Conference on Network Protocols (ICNP) , 2012.

- [IETFJ] Bonaventure, O. and S. Seo, "Multipath TCP Deployments," IETF Journal, Vol. 12, Issue 2 , November 2016.
- [IMC11] Honda, M., Nishida, Y., Raiciu, C., Greenhalgh, A., Handley, M., and H. Tokuda, "Is it still possible to extend TCP?", Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference (IMC '11) , 2011, <<http://doi.acm.org/10.1145/2068816.2068834>>.
- [IMC13a] Detal, G., Hesmans, B., Bonaventure, O., Vanaubel, Y., and B. Donnet, "Revealing Middlebox Interference with Tracebox", Proceedings of the 2013 ACM SIGCOMM conference on Internet measurement conference , 2013, <<http://inl.info.ucl.ac.be/publications/revealing-middlebox-interference-tracebox>>.
- [IMC13b] Chen, Y., Lim, Y., Gibbens, R., Nahum, E., Khalili, R., and D. Towsley, "A measurement-based study of MultiPath TCP performance over wireless network", Proceedings of the 2013 conference on Internet measurement conference (IMC '13) , n.d., <<http://doi.acm.org/10.1145/2504730.2504751>>.
- [IMC13c] Pelsser, C., Cittadini, L., Vissicchio, S., and R. Bush, "From Paris to Tokyo on the suitability of ping to measure latency", Proceedings of the 2013 conference on Internet measurement conference (IMC '13) , 2013, <<http://doi.acm.org/10.1145/2504730.2504765>>.
- [INFOCOM14] Lim, Y., Chen, Y., Nahum, E., Towsley, D., and K. Lee, "Cross-Layer Path Management in Multi-path Transport Protocol for Mobile Devices", IEEE INFOCOM'14 , 2014.
- [IOS7] Apple, ., "Multipath TCP Support in iOS 7", January 2014, <<http://support.apple.com/kb/HT5977>>.
- [KT] Seo, S., "KT's GiGA LTE", July 2015, <<https://www.ietf.org/proceedings/93/slides/slides-93-mptcp-3.pdf>>.
- [MBTest] Hesmans, B., "MBTest", 2013, <<https://bitbucket.org/bhesmans/mbtest>>.
- [MPTCPBIB] Bonaventure, O., "Multipath TCP - An annotated bibliography", Technical report , April 2015, <<https://github.com/obonaventure/mptcp-bib>>.
- [Mobicom15]

De Coninck, Q., Baerts, M., Hesmans, B., and O. Bonaventure, "Poster - Evaluating Android Applications with Multipath TCP", Mobicom 2015 (Poster) , September 2015.

[MultipathTCP-Linux]

Paasch, C., Barre, S., and . et al, "Multipath TCP implementation in the Linux kernel", n.d., <<http://www.multipath-tcp.org>>.

[NSDI11] Wischik, D., Raiciu, C., Greenhalgh, A., and M. Handley, "Design, implementation and evaluation of congestion control for Multipath TCP", In Proceedings of the 8th USENIX conference on Networked systems design and implementation (NSDI11) , 2011.

[NSDI12] Raiciu, C., Paasch, C., Barre, S., Ford, A., Honda, M., Duchene, F., Bonaventure, O., and M. Handley, "How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP", USENIX Symposium of Networked Systems Design and Implementation (NSDI12) , April 2012, <<http://inl.info.ucl.ac.be/publications/how-hard-can-it-be-designing-and-implementing-deployable-multipath-tcp>>.

[PAM2016] De Coninck, Q., Baerts, M., Hesmans, B., and O. Bonaventure, "A First Analysis of Multipath TCP on Smartphones", 17th International Passive and Active Measurements Conference (PAM2016) , March 2016, <<http://inl.info.ucl.ac.be/publications/first-analysis-multipath-tcp-smartphones>>.

[PAMS2014]

Arzani, B., Gurney, A., Cheng, S., Guerin, R., and B. Loo, "Impact of Path Selection and Scheduling Policies on MPTCP Performance", PAMS2014 , 2014.

[PT14] Pearce, C. and P. Thomas, "Multipath TCP Breaking Today's Networks with Tomorrow's Protocols", Proc. Blackhat Briefings , 2014, <<http://www.blackhat.com/docs/us-14/materials/us-14-Pearce-Multipath-TCP-Breaking-Todays-Networks-With-Tomorrows-Protocols-WP.pdf>>.

[PZ15] Pearce, C. and S. Zeadally, "Ancillary Impacts of Multipath TCP on Current and Future Network Security", IEEE Internet Computing, vol. 19, no. 5, pp. 58-65 , 2015.

- [PaaschPhD] Paasch, C., "Improving Multipath TCP", Ph.D. Thesis , November 2014, <<http://inl.info.ucl.ac.be/publications/improving-multipath-tcp>>.
- [Presto08] Greenberg, A., Lahiri, P., Maltz, D., Parveen, P., and S. Sengupta, "Towards a Next Generation Data Center Architecture - Scalability and Commoditization", ACM PRESTO 2008 , August 2008, <<http://dl.acm.org/citation.cfm?id=1397732>>.
- [RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers", RFC 1812, DOI 10.17487/RFC1812, June 1995, <<http://www.rfc-editor.org/info/rfc1812>>.
- [RFC1928] Leech, M., Ganis, M., Lee, Y., Kuris, R., Koblas, D., and L. Jones, "SOCKS Protocol Version 5", RFC 1928, DOI 10.17487/RFC1928, March 1996, <<http://www.rfc-editor.org/info/rfc1928>>.
- [RFC2992] Hopps, C., "Analysis of an Equal-Cost Multi-Path Algorithm", RFC 2992, DOI 10.17487/RFC2992, November 2000, <<http://www.rfc-editor.org/info/rfc2992>>.
- [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>>.
- [RFC6356] Raiciu, C., Handley, M., and D. Wischik, "Coupled Congestion Control for Multipath Transport Protocols", RFC 6356, DOI 10.17487/RFC6356, October 2011, <<http://www.rfc-editor.org/info/rfc6356>>.
- [RFC7430] Bagnulo, M., Paasch, C., Gont, F., Bonaventure, O., and C. Raiciu, "Analysis of Residual Threats and Possible Fixes for Multipath TCP (MPTCP)", RFC 7430, DOI 10.17487/RFC7430, July 2015, <<http://www.rfc-editor.org/info/rfc7430>>.
- [RFC7871] Contavalli, C., van der Gaast, W., Lawrence, D., and W. Kumari, "Client Subnet in DNS Queries", RFC 7871, DOI 10.17487/RFC7871, May 2016, <<http://www.rfc-editor.org/info/rfc7871>>.
- [SIGCOMM11] Raiciu, C., Barre, S., Pluntke, C., Greenhalgh, A., Wischik, D., and M. Handley, "Improving datacenter

performance and robustness with multipath TCP",
Proceedings of the ACM SIGCOMM 2011 conference , n.d.,
<<http://doi.acm.org/10.1145/2018436.2018467>>.

[SOCKET] Hesmans, B. and O. Bonaventure, "An Enhanced Socket API
for Multipath TCP", Proceedings of the 2016 Applied
Networking Research Workshop , July 2016,
<<http://doi.acm.org/10.1145/2959424.2959433>>.

[StrangeMbox]
Bonaventure, O., "Multipath TCP through a strange
middlebox", Blog post , January 2015, <[http://
blog.multipath-tcp.org/blog/html/2015/01/30/
multipath_tcp_through_a_strange_middlebox.html](http://blog.multipath-tcp.org/blog/html/2015/01/30/multipath_tcp_through_a_strange_middlebox.html)>.

[TMA2015] Hesmans, B., Tran Viet, H., Sadre, R., and O. Bonaventure,
"A First Look at Real Multipath TCP Traffic", Traffic
Monitoring and Analysis , 2015, <[http://
inl.info.ucl.ac.be/publications/
first-look-real-multipath-tcp-traffic](http://inl.info.ucl.ac.be/publications/first-look-real-multipath-tcp-traffic)>.

[TR-348] Broadband Forum, ., "TR 348 - Hybrid Access Broadband
Network Architecture", July 2016, <[https://
www.broadband-forum.org/technical/download/TR-348.pdf](https://www.broadband-forum.org/technical/download/TR-348.pdf)>.

[ietf88] Stewart, L., "IETF'88 Meeting minutes of the MPTCP working
group", n.d., <[http://tools.ietf.org/wg/mptcp/
minutes?item=minutes-88-mptcp.html](http://tools.ietf.org/wg/mptcp/minutes?item=minutes-88-mptcp.html)>.

[tracebox]
Detal, G. and O. Tilmans, "tracebox", 2013,
<<http://www.tracebox.org>>.

Appendix A. Changelog

This section should be removed before final publication

- o initial version : September 16th, 2014 : Added section Section 3.8 that discusses some performance problems that appeared with the Linux implementation when using subflows having different MSS values
- o update with a description of the middlebox that replaces an unknown TCP option with EOL [StrangeMbox]
- o version ietf-02 : July 2015, answer to last call comments
 - * Reorganised text to better separate use cases and operational experience
 - * New use case on Multipath TCP proxies in Section 2.3
 - * Added some text on middleboxes in Section 3.1
 - * Removed the discussion on SDN
 - * Restructured text and improved writing in some parts
- o version ietf-03 : September 2015, answer to comments from Phil Eardley
 - * Improved introduction
 - * Added details about using SOCKS and Korea Telecom's use-case in Section 2.3.
 - * Added issue around clients caching DNS-results in Section 3.9
 - * Explained issue of MPTCP with stateless webserver Section 3.11
 - * Added description of MPTCP's use behind layer-4 load balancers Section 3.12
 - * Restructured text and improved writing in some parts
- o version ietf-04 : April 2016, answer to last comments
 - * Updated text on measurements with smartphones
 - * Updated conclusion

- o version ietf-06 : August 2016, answer to AD's review
 - * removed some expired drafts
 - * removed conclusion
- o version ietf-07 : September 2016, answer to IESG comments
 - * small nits
 - * added more information in the security considerations as suggested by Stephen Farrel

Authors' Addresses

Olivier Bonaventure
UCLouvain

Email: Olivier.Bonaventure@uclouvain.be

Christoph Paasch
Apple, Inc.

Email: cpaasch@apple.com

Gregory Detal
Tessares

Email: Gregory.Detal@tessares.net

Internet Engineering Task Force
Internet-Draft
Obsoletes: 6824 (if approved)
Intended status: Standards Track
Expires: December 10, 2019

A. Ford
Pexip
C. Raiciu
U. Politechnica of Bucharest
M. Handley
U. College London
O. Bonaventure
U. catholique de Louvain
C. Paasch
Apple, Inc.
June 8, 2019

TCP Extensions for Multipath Operation with Multiple Addresses
draft-ietf-mptcp-rfc6824bis-18

Abstract

TCP/IP communication is currently restricted to a single path per connection, yet multiple paths often exist between peers. The simultaneous use of these multiple paths for a TCP/IP session would improve resource usage within the network and, thus, improve user experience through higher throughput and improved resilience to network failure.

Multipath TCP provides the ability to simultaneously use multiple paths between peers. This document presents a set of extensions to traditional TCP to support multipath operation. The protocol offers the same type of service to applications as TCP (i.e., reliable bytestream), and it provides the components necessary to establish and use multiple TCP flows across potentially disjoint paths.

This document specifies v1 of Multipath TCP, obsoleting v0 as specified in RFC6824, through clarifications and modifications primarily driven by deployment experience.

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 December 10, 2019.

Copyright Notice

Copyright (c) 2019 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	3
1.1.	Design Assumptions	4
1.2.	Multipath TCP in the Networking Stack	5
1.3.	Terminology	6
1.4.	MPTCP Concept	7
1.5.	Requirements Language	8
2.	Operation Overview	8
2.1.	Initiating an MPTCP Connection	9
2.2.	Associating a New Subflow with an Existing MPTCP Connection	10
2.3.	Informing the Other Host about Another Potential Address	11
2.4.	Data Transfer Using MPTCP	12
2.5.	Requesting a Change in a Path's Priority	13
2.6.	Closing an MPTCP Connection	13
2.7.	Notable Features	14
3.	MPTCP Protocol	15
3.1.	Connection Initiation	16
3.2.	Starting a New Subflow	23
3.3.	General MPTCP Operation	28
3.3.1.	Data Sequence Mapping	30
3.3.2.	Data Acknowledgments	33
3.3.3.	Closing a Connection	34
3.3.4.	Receiver Considerations	36
3.3.5.	Sender Considerations	37
3.3.6.	Reliability and Retransmissions	38
3.3.7.	Congestion Control Considerations	39

3.3.8. Subflow Policy	39
3.4. Address Knowledge Exchange (Path Management)	40
3.4.1. Address Advertisement	42
3.4.2. Remove Address	45
3.5. Fast Close	46
3.6. Subflow Reset	48
3.7. Fallback	49
3.8. Error Handling	53
3.9. Heuristics	53
3.9.1. Port Usage	54
3.9.2. Delayed Subflow Start and Subflow Symmetry	54
3.9.3. Failure Handling	55
4. Semantic Issues	56
5. Security Considerations	57
6. Interactions with Middleboxes	60
7. Acknowledgments	63
8. IANA Considerations	64
8.1. MPTCP Option Subtypes	64
8.2. MPTCP Handshake Algorithms	65
8.3. MP_TCPRST Reason Codes	66
9. References	67
9.1. Normative References	67
9.2. Informative References	68
Appendix A. Notes on Use of TCP Options	71
Appendix B. TCP Fast Open and MPTCP	72
B.1. TFO cookie request with MPTCP	72
B.2. Data sequence mapping under TFO	73
B.3. Connection establishment examples	74
Appendix C. Control Blocks	76
C.1. MPTCP Control Block	76
C.1.1. Authentication and Metadata	76
C.1.2. Sending Side	77
C.1.3. Receiving Side	77
C.2. TCP Control Blocks	77
C.2.1. Sending Side	78
C.2.2. Receiving Side	78
Appendix D. Finite State Machine	78
Appendix E. Changes from RFC6824	79
Authors' Addresses	81

1. Introduction

Multipath TCP (MPTCP) is a set of extensions to regular TCP [RFC0793] to provide a Multipath TCP [RFC6182] service, which enables a transport connection to operate across multiple paths simultaneously. This document presents the protocol changes required to add multipath capability to TCP; specifically, those for signaling and setting up multiple paths ("subflows"), managing these subflows, reassembly of

data, and termination of sessions. This is not the only information required to create a Multipath TCP implementation, however. This document is complemented by three others:

- o Architecture [RFC6182], which explains the motivations behind Multipath TCP, contains a discussion of high-level design decisions on which this design is based, and an explanation of a functional separation through which an extensible MPTCP implementation can be developed.
- o Congestion control [RFC6356] presents a safe congestion control algorithm for coupling the behavior of the multiple paths in order to "do no harm" to other network users.
- o Application considerations [RFC6897] discusses what impact MPTCP will have on applications, what applications will want to do with MPTCP, and as a consequence of these factors, what API extensions an MPTCP implementation should present.

This document is an update to, and obsoletes, the v0 specification of Multipath TCP (RFC6824). This document specifies MPTCP v1, which is not backward compatible with MPTCP v0. This document additionally defines version negotiation procedures for implementations that support both versions.

1.1. Design Assumptions

In order to limit the potentially huge design space, the mptcp working group imposed two key constraints on the Multipath TCP design presented in this document:

- o It must be backwards-compatible with current, regular TCP, to increase its chances of deployment.
- o It can be assumed that one or both hosts are multihomed and multiaddressed.

To simplify the design, we assume that the presence of multiple addresses at a host is sufficient to indicate the existence of multiple paths. These paths need not be entirely disjoint: they may share one or many routers between them. Even in such a situation, making use of multiple paths is beneficial, improving resource utilization and resilience to a subset of node failures. The congestion control algorithms defined in [RFC6356] ensure this does not act detrimentally. Furthermore, there may be some scenarios where different TCP ports on a single host can provide disjoint paths (such as through certain Equal-Cost Multipath (ECMP) implementations

[RFC2992]), and so the MPTCP design also supports the use of ports in path identifiers.

There are three aspects to the backwards-compatibility listed above (discussed in more detail in [RFC6182]):

External Constraints: The protocol must function through the vast majority of existing middleboxes such as NATs, firewalls, and proxies, and as such must resemble existing TCP as far as possible on the wire. Furthermore, the protocol must not assume the segments it sends on the wire arrive unmodified at the destination: they may be split or coalesced; TCP options may be removed or duplicated.

Application Constraints: The protocol must be usable with no change to existing applications that use the common TCP API (although it is reasonable that not all features would be available to such legacy applications). Furthermore, the protocol must provide the same service model as regular TCP to the application.

Fallback: The protocol should be able to fall back to standard TCP with no interference from the user, to be able to communicate with legacy hosts.

The complementary application considerations document [RFC6897] discusses the necessary features of an API to provide backwards-compatibility, as well as API extensions to convey the behavior of MPTCP at a level of control and information equivalent to that available with regular, single-path TCP.

Further discussion of the design constraints and associated design decisions are given in the MPTCP Architecture document [RFC6182] and in [howhard].

1.2. Multipath TCP in the Networking Stack

MPTCP operates at the transport layer and aims to be transparent to both higher and lower layers. It is a set of additional features on top of standard TCP; Figure 1 illustrates this layering. MPTCP is designed to be usable by legacy applications with no changes; detailed discussion of its interactions with applications is given in [RFC6897].

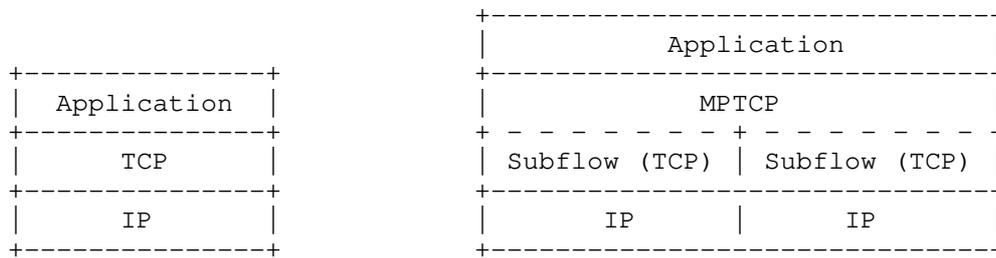


Figure 1: Comparison of Standard TCP and MPTCP Protocol Stacks

1.3. Terminology

This document makes use of a number of terms that are either MPTCP-specific or have defined meaning in the context of MPTCP, as follows:

Path: A sequence of links between a sender and a receiver, defined in this context by a 4-tuple of source and destination address/port pairs.

Subflow: A flow of TCP segments operating over an individual path, which forms part of a larger MPTCP connection. A subflow is started and terminated similar to a regular TCP connection.

(MPTCP) Connection: A set of one or more subflows, over which an application can communicate between two hosts. There is a one-to-one mapping between a connection and an application socket.

Data-level: The payload data is nominally transferred over a connection, which in turn is transported over subflows. Thus, the term "data-level" is synonymous with "connection level", in contrast to "subflow-level", which refers to properties of an individual subflow.

Token: A locally unique identifier given to a multipath connection by a host. May also be referred to as a "Connection ID".

Host: An end host operating an MPTCP implementation, and either initiating or accepting an MPTCP connection.

In addition to these terms, note that MPTCP's interpretation of, and effect on, regular single-path TCP semantics are discussed in Section 4.

1.4. MPTCP Concept

This section provides a high-level summary of normal operation of MPTCP, and is illustrated by the scenario shown in Figure 2. A detailed description of operation is given in Section 3.

- o To a non-MPTCP-aware application, MPTCP will behave the same as normal TCP. Extended APIs could provide additional control to MPTCP-aware applications [RFC6897]. An application begins by opening a TCP socket in the normal way. MPTCP signaling and operation are handled by the MPTCP implementation.
- o An MPTCP connection begins similarly to a regular TCP connection. This is illustrated in Figure 2 where an MPTCP connection is established between addresses A1 and B1 on Hosts A and B, respectively.
- o If extra paths are available, additional TCP sessions (termed MPTCP "subflows") are created on these paths, and are combined with the existing session, which continues to appear as a single connection to the applications at both ends. The creation of the additional TCP session is illustrated between Address A2 on Host A and Address B1 on Host B.
- o MPTCP identifies multiple paths by the presence of multiple addresses at hosts. Combinations of these multiple addresses equate to the additional paths. In the example, other potential paths that could be set up are A1<->B2 and A2<->B2. Although this additional session is shown as being initiated from A2, it could equally have been initiated from B1 or B2.
- o The discovery and setup of additional subflows will be achieved through a path management method; this document describes a mechanism by which a host can initiate new subflows by using its own additional addresses, or by signaling its available addresses to the other host.
- o MPTCP adds connection-level sequence numbers to allow the reassembly of segments arriving on multiple subflows with differing network delays.
- o Subflows are terminated as regular TCP connections, with a four-way FIN handshake. The MPTCP connection is terminated by a connection-level FIN.

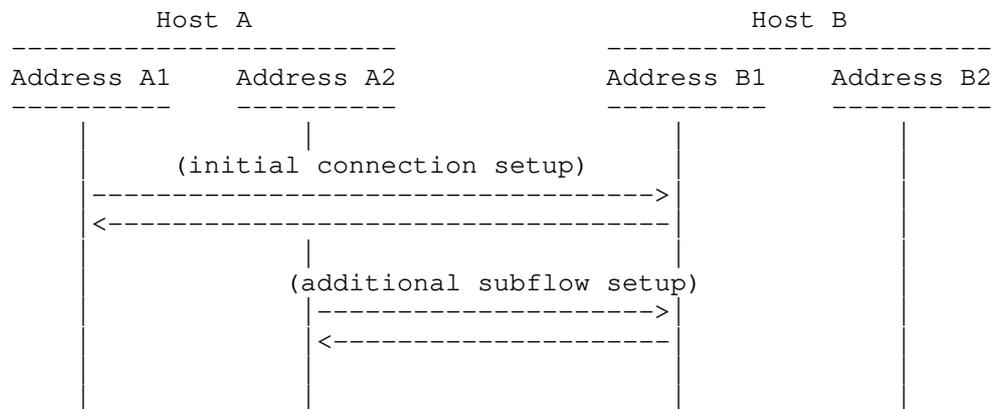


Figure 2: Example MPTCP Usage Scenario

1.5. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Operation Overview

This section presents a single description of common MPTCP operation, with reference to the protocol operation. This is a high-level overview of the key functions; the full specification follows in Section 3. Extensibility and negotiated features are not discussed here. Considerable reference is made to symbolic names of MPTCP options throughout this section -- these are subtypes of the IANA-assigned MPTCP option (see Section 8), and their formats are defined in the detailed protocol specification that follows in Section 3.

A Multipath TCP connection provides a bidirectional bytestream between two hosts communicating like normal TCP and, thus, does not require any change to the applications. However, Multipath TCP enables the hosts to use different paths with different IP addresses to exchange packets belonging to the MPTCP connection. A Multipath TCP connection appears like a normal TCP connection to an application. However, to the network layer, each MPTCP subflow looks like a regular TCP flow whose segments carry a new TCP option type. Multipath TCP manages the creation, removal, and utilization of these subflows to send data. The number of subflows that are managed within a Multipath TCP connection is not fixed and it can fluctuate during the lifetime of the Multipath TCP connection.

All MPTCP operations are signaled with a TCP option -- a single numerical type for MPTCP, with "sub-types" for each MPTCP message. What follows is a summary of the purpose and rationale of these messages.

2.1. Initiating an MPTCP Connection

This is the same signaling as for initiating a normal TCP connection, but the SYN, SYN/ACK, and initial ACK (and data) packets also carry the MP_CAPABLE option. This option has a variable length and serves multiple purposes. Firstly, it verifies whether the remote host supports Multipath TCP; secondly, this option allows the hosts to exchange some information to authenticate the establishment of additional subflows. Further details are given in Section 3.1.

```

Host A                               Host B
-----                               -----
MP_CAPABLE                            ->
[flags]

                                     <-
                                     MP_CAPABLE
                                     [B's key, flags]

ACK + MP_CAPABLE (+ data) ->
[A's key, B's key, flags, (data-level details)]

```

Retransmission of the ACK + MP_CAPABLE can occur if it is not known if it has been received. The following diagrams show all possible exchanges for the initial subflow setup to ensure this reliability.

```

Host A (with data to send immediately)  Host B
-----
MP_CAPABLE                               ->
[flags]
                                         <-
                                         MP_CAPABLE
                                         [B's key, flags]

ACK + MP_CAPABLE + data                  ->
[A's key, B's key, flags, data-level details]

Host A (with data to send later)        Host B
-----
MP_CAPABLE                               ->
[flags]
                                         <-
                                         MP_CAPABLE
                                         [B's key, flags]

ACK + MP_CAPABLE                         ->
[A's key, B's key, flags]

ACK + MP_CAPABLE + data                  ->
[A's key, B's key, flags, data-level details]

Host A                                    Host B (sending first)
-----
MP_CAPABLE                               ->
[flags]
                                         <-
                                         MP_CAPABLE
                                         [B's key, flags]

ACK + MP_CAPABLE                         ->
[A's key, B's key, flags]

                                         <-
                                         ACK + DSS + data
                                         [data-level details]

```

2.2. Associating a New Subflow with an Existing MPTCP Connection

The exchange of keys in the MP_CAPABLE handshake provides material that can be used to authenticate the endpoints when new subflows will be set up. Additional subflows begin in the same way as initiating a normal TCP connection, but the SYN, SYN/ACK, and ACK packets also carry the MP_JOIN option.

Host A initiates a new subflow between one of its addresses and one of Host B's addresses. The token -- generated from the key -- is used to identify which MPTCP connection it is joining, and the HMAC is used for authentication. The Hash-based Message Authentication Code (HMAC) uses the keys exchanged in the MP_CAPABLE handshake, and

the random numbers (nonces) exchanged in these MP_JOIN options. MP_JOIN also contains flags and an Address ID that can be used to refer to the source address without the sender needing to know if it has been changed by a NAT. Further details are in Section 3.2.

```

Host A                               Host B
-----                               -----
MP_JOIN                               ->
[B's token, A's nonce,
 A's Address ID, flags]
<-
ACK + MP_JOIN                         ->
[A's HMAC]
<-                                     ACK
MP_JOIN
[B's HMAC, B's nonce,
 B's Address ID, flags]
```

2.3. Informing the Other Host about Another Potential Address

The set of IP addresses associated to a multihomed host may change during the lifetime of an MPTCP connection. MPTCP supports the addition and removal of addresses on a host both implicitly and explicitly. If Host A has established a subflow starting at address/port pair IP#-A1 and wants to open a second subflow starting at address/port pair IP#-A2, it simply initiates the establishment of the subflow as explained above. The remote host will then be implicitly informed about the new address.

In some circumstances, a host may want to advertise to the remote host the availability of an address without establishing a new subflow, for example, when a NAT prevents setup in one direction. In the example below, Host A informs Host B about its alternative IP address/port pair (IP#-A2). Host B may later send an MP_JOIN to this new address. The ADD_ADDR option contains a HMAC to authenticate the address as having been sent from the originator of the connection. The receiver of this option echoes it back to the client to indicate successful receipt. Further details are in Section 3.4.1.

```

Host A                                     Host B
-----                                     -----
ADD_ADDR                                  ->
[Echo-flag=0,
 IP#-A2,
 IP#-A2's Address ID,
 HMAC of IP#-A2]

<-
ADD_ADDR
[Echo-flag=1,
 IP#-A2,
 IP#-A2's Address ID,
 HMAC of IP#-A2]

```

There is a corresponding signal for address removal, making use of the Address ID that is signaled in the add address handshake. Further details in Section 3.4.2.

```

Host A                                     Host B
-----                                     -----
REMOVE_ADDR                               ->
[IP#-A2's Address ID]

```

2.4. Data Transfer Using MPTCP

To ensure reliable, in-order delivery of data over subflows that may appear and disappear at any time, MPTCP uses a 64-bit data sequence number (DSN) to number all data sent over the MPTCP connection. Each subflow has its own 32-bit sequence number space, utilising the regular TCP sequence number header, and an MPTCP option maps the subflow sequence space to the data sequence space. In this way, data can be retransmitted on different subflows (mapped to the same DSN) in the event of failure.

The Data Sequence Signal (DSS) carries the Data Sequence Mapping. The Data Sequence Mapping consists of the subflow sequence number, data sequence number, and length for which this mapping is valid. This option can also carry a connection-level acknowledgment (the "Data ACK") for the received DSN.

With MPTCP, all subflows share the same receive buffer and advertise the same receive window. There are two levels of acknowledgment in MPTCP. Regular TCP acknowledgments are used on each subflow to acknowledge the reception of the segments sent over the subflow independently of their DSN. In addition, there are connection-level acknowledgments for the data sequence space. These acknowledgments track the advancement of the bytestream and slide the receiving window.

Further details are in Section 3.3.

```

Host A                               Host B
-----                               -----
DSS                                   ->
[Data Sequence Mapping]
[Data ACK]
[Checksum]

```

2.5. Requesting a Change in a Path's Priority

Hosts can indicate at initial subflow setup whether they wish the subflow to be used as a regular or backup path -- a backup path only being used if there are no regular paths available. During a connection, Host A can request a change in the priority of a subflow through the MP_PRIO signal to Host B. Further details are in Section 3.3.8.

```

Host A                               Host B
-----                               -----
MP_PRIO                              ->

```

2.6. Closing an MPTCP Connection

When a host wants to close an existing subflow, but not the whole connection, it can initiate a regular TCP FIN/ACK exchange.

When Host A wants to inform Host B that it has no more data to send, it signals this "Data FIN" as part of the Data Sequence Signal (see above). It has the same semantics and behavior as a regular TCP FIN, but at the connection level. Once all the data on the MPTCP connection has been successfully received, then this message is acknowledged at the connection level with a Data ACK. Further details are in Section 3.3.3.

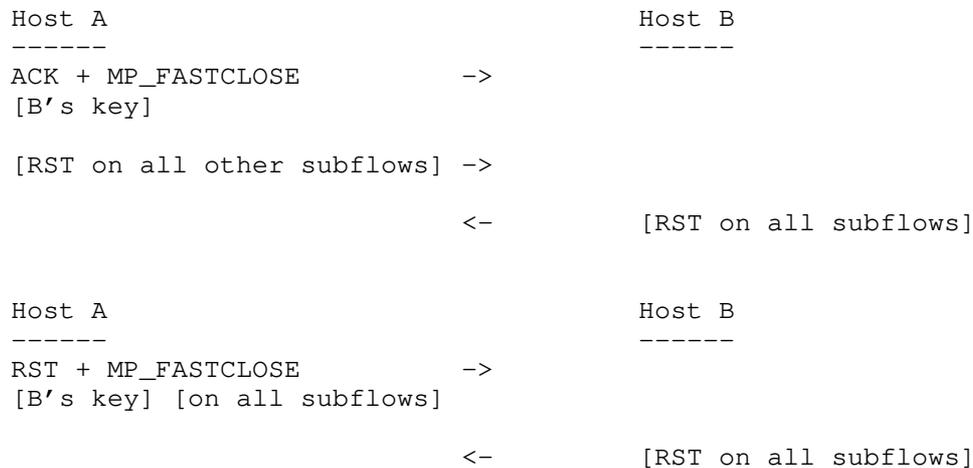
```

Host A                               Host B
-----                               -----
DSS                                   ->
[Data FIN]
<-
DSS
[Data ACK]

```

There is an additional method of connection closure, referred to as "Fast Close", which is analogous to closing a single-path TCP connection with a RST signal. The MP_FASTCLOSE signal is used to indicate to the peer that the connection will be abruptly closed and no data will be accepted anymore. This can be used on an ACK (ensuring reliability of the signal), or a RST (which is not). Both

examples are shown in the following diagrams. Further details are in Section 3.5.



2.7. Notable Features

It is worth highlighting that MPTCP's signaling has been designed with several key requirements in mind:

- o To cope with NATs on the path, addresses are referred to by Address IDs, in case the IP packet's source address gets changed by a NAT. Setting up a new TCP flow is not possible if the receiver of the SYN is behind a NAT; to allow subflows to be created when either end is behind a NAT, MPTCP uses the ADD_ADDR message.
- o MPTCP falls back to ordinary TCP if MPTCP operation is not possible, for example, if one host is not MPTCP capable or if a middlebox alters the payload. This is discussed in Section 3.7.
- o To address the threats identified in [RFC6181], the following steps are taken: keys are sent in the clear in the MP_CAPABLE messages; MP_JOIN messages are secured with HMAC-SHA256 ([RFC2104], [RFC6234]) using those keys; and standard TCP validity checks are made on the other messages (ensuring sequence numbers are in-window [RFC5961]). Residual threats to MPTCP v0 were identified in [RFC7430], and those affecting the protocol (i.e. modification to ADD_ADDR) have been incorporated in this document. Further discussion of security can be found in Section 5.

3. MPTCP Protocol

This section describes the operation of the MPTCP protocol, and is subdivided into sections for each key part of the protocol operation.

All MPTCP operations are signaled using optional TCP header fields. A single TCP option number ("Kind") has been assigned by IANA for MPTCP (see Section 8), and then individual messages will be determined by a "subtype", the values of which are also stored in an IANA registry (and are also listed in Section 8). As with all TCP options, the Length field is specified in bytes, and includes the 2 bytes of Kind and Length.

Throughout this document, when reference is made to an MPTCP option by symbolic name, such as "MP_CAPABLE", this refers to a TCP option with the single MPTCP option type, and with the subtype value of the symbolic name as defined in Section 8. This subtype is a 4-bit field -- the first 4 bits of the option payload, as shown in Figure 3. The MPTCP messages are defined in the following sections.

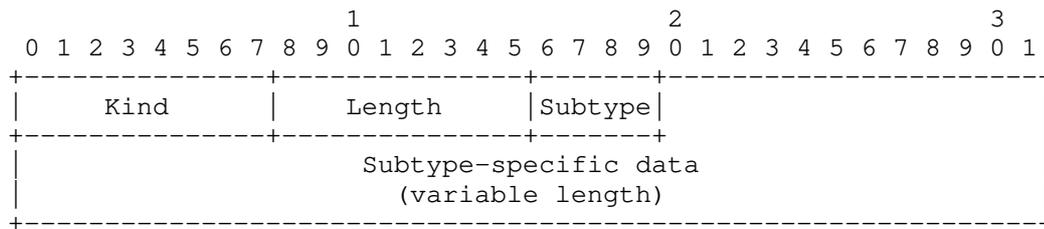


Figure 3: MPTCP Option Format

Those MPTCP options associated with subflow initiation are used on packets with the SYN flag set. Additionally, there is one MPTCP option for signaling metadata to ensure segmented data can be recombined for delivery to the application.

The remaining options, however, are signals that do not need to be on a specific packet, such as those for signaling additional addresses. Whilst an implementation may desire to send MPTCP options as soon as possible, it may not be possible to combine all desired options (both those for MPTCP and for regular TCP, such as SACK (selective acknowledgment) [RFC2018]) on a single packet. Therefore, an implementation may choose to send duplicate ACKs containing the additional signaling information. This changes the semantics of a duplicate ACK; these are usually only sent as a signal of a lost segment [RFC5681] in regular TCP. Therefore, an MPTCP implementation receiving a duplicate ACK that contains an MPTCP option MUST NOT treat it as a signal of congestion. Additionally, an MPTCP

implementation SHOULD NOT send more than two duplicate ACKs in a row for the purposes of sending MPTCP options alone, in order to ensure no middleboxes misinterpret this as a sign of congestion.

Furthermore, standard TCP validity checks (such as ensuring the sequence number and acknowledgment number are within window) MUST be undertaken before processing any MPTCP signals, as described in [RFC5961], and initial subflow sequence numbers SHOULD be generated according to the recommendations in [RFC6528].

3.1. Connection Initiation

Connection initiation begins with a SYN, SYN/ACK, ACK exchange on a single path. Each packet contains the Multipath Capable (MP_CAPABLE) MPTCP option (Figure 4). This option declares its sender is capable of performing Multipath TCP and wishes to do so on this particular connection.

The MP_CAPABLE exchange in this specification (v1) is different to that specified in v0. If a host supports multiple versions of MPTCP, the sender of the MP_CAPABLE option SHOULD signal the highest version number it supports. In return, in its MP_CAPABLE option, the receiver will signal the version number it wishes to use, which MUST be equal to or lower than the version number indicated in the initial MP_CAPABLE. There is a caveat though with respect to this version negotiation with old listeners that only support v0. A listener that supports v0 expects that the MP_CAPABLE option in the SYN-segment includes the initiator's key. If the initiator however already upgraded to v1, it won't include the key in the SYN-segment. Thus, the listener will ignore the MP_CAPABLE of this SYN-segment and reply with a SYN/ACK that does not include an MP_CAPABLE. The initiator MAY choose to immediately fall back to TCP or MAY choose to attempt a connection using MPTCP v0 (if the initiator supports v0), in order to discover whether the listener supports the earlier version of MPTCP. In general a MPTCP v0 connection is likely to be preferred to a TCP one, however in a particular deployment scenario it may be known that the listener is unlikely to support MPTCPv0 and so the initiator may prefer not to attempt a v0 connection. An initiator MAY cache information for a peer about what version of MPTCP it supports if any, and use this information for future connection attempts.

The MP_CAPABLE option is variable-length, with different fields included depending on which packet the option is used on. The full MP_CAPABLE option is shown in Figure 4.

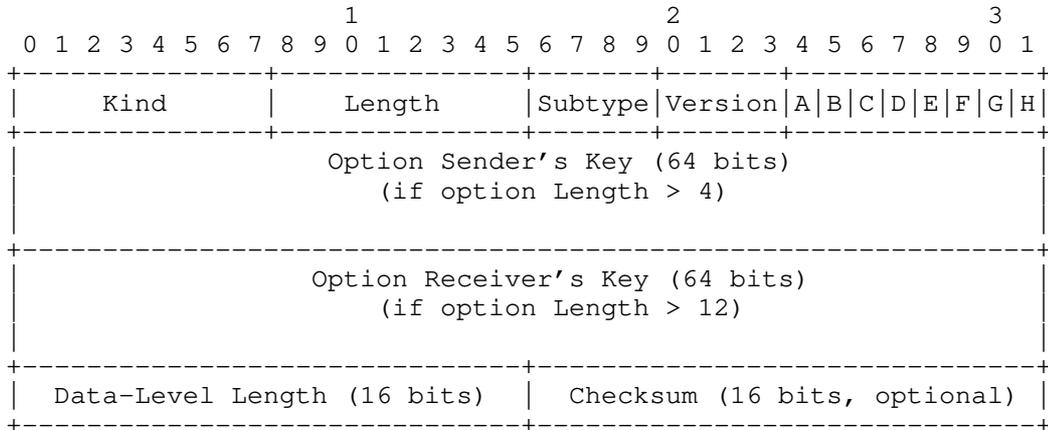


Figure 4: Multipath Capable (MP_CAPABLE) Option

The MP_CAPABLE option is carried on the SYN, SYN/ACK, and ACK packets that start the first subflow of an MPTCP connection, as well as the first packet that carries data, if the initiator wishes to send first. The data carried by each option is as follows, where A = initiator and B = listener.

- o SYN (A->B): only the first four octets (Length = 4).
- o SYN/ACK (B->A): B's Key for this connection (Length = 12).
- o ACK (no data) (A->B): A's Key followed by B's Key (Length = 20).
- o ACK (with first data) (A->B): A's Key followed by B's Key followed by Data-Level Length, and optional Checksum (Length = 22 or 24).

The contents of the option is determined by the SYN and ACK flags of the packet, along with the option's length field. For the diagram shown in Figure 4, "sender" and "receiver" refer to the sender or receiver of the TCP packet (which can be either host).

The initial SYN, containing just the MP_CAPABLE header, is used to define the version of MPTCP being requested, as well as exchanging flags to negotiate connection features, described later.

This option is used to declare the 64-bit keys that the end hosts have generated for this MPTCP connection. These keys are used to authenticate the addition of future subflows to this connection. This is the only time the key will be sent in clear on the wire (unless "fast close", Section 3.5, is used); all future subflows will identify the connection using a 32-bit "token". This token is a

cryptographic hash of this key. The algorithm for this process is dependent on the authentication algorithm selected; the method of selection is defined later in this section.

Upon reception of the initial SYN-segment, a stateful server generates a random key and replies with a SYN/ACK. The key's method of generation is implementation specific. The key MUST be hard to guess, and it MUST be unique for the sending host across all its current MPTCP connections. Recommendations for generating random numbers for use in keys are given in [RFC4086]. Connections will be indexed at each host by the token (a one-way hash of the key). Therefore, an implementation will require a mapping from each token to the corresponding connection, and in turn to the keys for the connection.

There is a risk that two different keys will hash to the same token. The risk of hash collisions is usually small, unless the host is handling many tens of thousands of connections. Therefore, an implementation SHOULD check its list of connection tokens to ensure there is no collision before sending its key, and if there is, then it should generate a new key. This would, however, be costly for a server with thousands of connections. The subflow handshake mechanism (Section 3.2) will ensure that new subflows only join the correct connection, however, through the cryptographic handshake, as well as checking the connection tokens in both directions, and ensuring sequence numbers are in-window. So in the worst case if there was a token collision, the new subflow would not succeed, but the MPTCP connection would continue to provide a regular TCP service.

Since key generation is implementation-specific, there is no requirement that they be simply random numbers. An implementation is free to exchange cryptographic material out-of-band and generate these keys from this, in order to provide additional mechanisms by which to verify the identity of the communicating entities. For example, an implementation could choose to link its MPTCP keys to those used in higher-layer TLS or SSH connections.

If the server behaves in a stateless manner, it has to generate its own key in a verifiable fashion. This verifiable way of generating the key can be done by using a hash of the 4-tuple, sequence number and a local secret (similar to what is done for the TCP-sequence number [RFC4987]). It will thus be able to verify whether it is indeed the originator of the key echoed back in the later MP_CAPABLE option. As for a stateful server, the tokens SHOULD be checked for uniqueness, however if uniqueness is not met, and there is no way to generate an alternative verifiable key, then the connection MUST fall back to using regular TCP by not sending a MP_CAPABLE in the SYN/ACK.

The ACK carries both A's key and B's key. This is the first time that A's key is seen on the wire, although it is expected that A will have generated a key locally before the initial SYN. The echoing of B's key allows B to operate statelessly, as described above. Therefore, A's key must be delivered reliably to B, and in order to do this, the transmission of this packet must be made reliable.

If B has data to send first, then the reliable delivery of the ACK+MP_CAPABLE can be inferred by the receipt of this data with a MPTCP Data Sequence Signal (DSS) option (Section 3.3). If, however, A wishes to send data first, it has two options to ensure the reliable delivery of the ACK+MP_CAPABLE. If it immediately has data to send, then the third ACK (with data) would also contain an MP_CAPABLE option with additional data parameters (the Data-Level Length and optional Checksum as shown in Figure 4). If A does not immediately have data to send, it MUST include the MP_CAPABLE on the third ACK, but without the additional data parameters. When A does have data to send, it must repeat the sending of the MP_CAPABLE option from the third ACK, with additional data parameters. This MP_CAPABLE option is in place of the DSS, and simply specifies the data-level length of the payload, and the checksum (if the use of checksums is negotiated). This is the minimal data required to establish a MPTCP connection - it allows validation of the payload, and given it is the first data, the Initial Data Sequence Number (IDSN) is also known (as it is generated from the key, as described below). Conveying the keys on the first data packet allows the TCP reliability mechanisms to ensure the packet is successfully delivered. The receiver will acknowledge this data at the connection level with a Data ACK, as if a DSS option has been received.

There could be situations where both A and B attempt to transmit initial data at the same time. For example, if A did not initially have data to send, but then needed to transmit data before it had received anything from B, it would use a MP_CAPABLE option with data parameters (since it would not know if the MP_CAPABLE on the ACK was received). In such a situation, B may also have transmitted data with a DSS option, but it had not yet been received at A. Therefore, B has received data with a MP_CAPABLE mapping after it has sent data with a DSS option. To ensure these situations can be handled, it follows that the data parameters in a MP_CAPABLE are semantically equivalent to those in a DSS option and can be used interchangeably. Similar situations could occur when the MP_CAPABLE with data is lost and retransmitted. Furthermore, in the case of TCP Segmentation Offloading, the MP_CAPABLE with data parameters may be duplicated across multiple packets, and implementations must also be able to cope with duplicate MP_CAPABLE mappings as well as duplicate DSS mappings.

Additionally, the MP_CAPABLE exchange allows the safe passage of MPTCP options on SYN packets to be determined. If any of these options are dropped, MPTCP will gracefully fall back to regular single-path TCP, as documented in Section 3.7. If at any point in the handshake either party thinks the MPTCP negotiation is compromised, for example by a middlebox corrupting the TCP options, or unexpected ACK numbers being present, the host MUST stop using MPTCP and no longer include MPTCP options in future TCP packets. The other host will then also fall back to regular TCP using the fall back mechanism. Note that new subflows MUST NOT be established (using the process documented in Section 3.2) until a Data Sequence Signal (DSS) option has been successfully received across the path (as documented in Section 3.3).

Like all MPTCP options, the MP_CAPABLE option starts with the Kind and Length to specify the TCP-option kind and its length. Followed by that is the MP_CAPABLE option. The first 4 bits of the first octet in the MP_CAPABLE option (Figure 4) define the MPTCP option subtype (see Section 8; for MP_CAPABLE, this is 0x0), and the remaining 4 bits of this octet specify the MPTCP version in use (for this specification, this is 1).

The second octet is reserved for flags, allocated as follows:

- A: The leftmost bit, labeled "A", SHOULD be set to 1 to indicate "Checksum Required", unless the system administrator has decided that checksums are not required (for example, if the environment is controlled and no middleboxes exist that might adjust the payload).
- B: The second bit, labeled "B", is an extensibility flag, and MUST be set to 0 for current implementations. This will be used for an extensibility mechanism in a future specification, and the impact of this flag will be defined at a later date. It is expected, but not mandated, that this flag would be used as part of an alternative security mechanism that does not require a full version upgrade of the protocol, but does require redefining some elements of the handshake. If receiving a message with the 'B' flag set to 1, and this is not understood, then the MP_CAPABLE in this SYN MUST be silently ignored, which triggers a fallback to regular TCP; the sender is expected to retry with a format compatible with this legacy specification. Note that the length of the MP_CAPABLE option, and the meanings of bits "D" through "H", may be altered by setting B=1.
- C: The third bit, labeled "C", is set to "1" to indicate that the sender of this option will not accept additional MPTCP subflows to the source address and port, and therefore the receiver MUST NOT

try to open any additional subflows towards this address and port. This is an efficiency improvement for situations where the sender knows a restriction is in place, for example if the sender is behind a strict NAT, or operating behind a legacy Layer 4 load balancer.

D through H: The remaining bits, labeled "D" through "H", are used for crypto algorithm negotiation. In this specification only the rightmost bit, labeled "H", is assigned. Bit "H" indicates the use of HMAC-SHA256 (as defined in Section 3.2). An implementation that only supports this method MUST set bit "H" to 1, and bits "D" through "G" to 0.

A crypto algorithm MUST be specified. If flag bits D through H are all 0, the MP_CAPABLE option MUST be treated as invalid and ignored (that is, it must be treated as a regular TCP handshake).

The selection of the authentication algorithm also impacts the algorithm used to generate the token and the Initial Data Sequence Number (IDSN). In this specification, with only the SHA-256 algorithm (bit "H") specified and selected, the token MUST be a truncated (most significant 32 bits) SHA-256 hash ([RFC6234]) of the key. A different, 64-bit truncation (the least significant 64 bits) of the SHA-256 hash of the key MUST be used as the IDSN. Note that the key MUST be hashed in network byte order. Also note that the "least significant" bits MUST be the rightmost bits of the SHA-256 digest, as per [RFC6234]. Future specifications of the use of the crypto bits may choose to specify different algorithms for token and IDSN generation.

Both the crypto and checksum bits negotiate capabilities in similar ways. For the Checksum Required bit (labeled "A"), if either host requires the use of checksums, checksums MUST be used. In other words, the only way for checksums not to be used is if both hosts in their SYNs set A=0. This decision is confirmed by the setting of the "A" bit in the third packet (the ACK) of the handshake. For example, if the initiator sets A=0 in the SYN, but the responder sets A=1 in the SYN/ACK, checksums MUST be used in both directions, and the initiator will set A=1 in the ACK. The decision whether to use checksums will be stored by an implementation in a per-connection binary state variable. If A=1 is received by a host that does not want to use checksums, it MUST fall back to regular TCP by ignoring the MP_CAPABLE option as if it was invalid.

For crypto negotiation, the responder has the choice. The initiator creates a proposal setting a bit for each algorithm it supports to 1 (in this version of the specification, there is only one proposal, so bit "H" will be always set to 1). The responder responds with only 1

bit set -- this is the chosen algorithm. The rationale for this behavior is that the responder will typically be a server with potentially many thousands of connections, so it may wish to choose an algorithm with minimal computational complexity, depending on the load. If a responder does not support (or does not want to support) any of the initiator's proposals, it MUST respond without an MP_CAPABLE option, thus forcing a fallback to regular TCP.

The MP_CAPABLE option is only used in the first subflow of a connection, in order to identify the connection; all following subflows will use the "Join" option (see Section 3.2) to join the existing connection.

If a SYN contains an MP_CAPABLE option but the SYN/ACK does not, it is assumed that sender of the SYN/ACK is not multipath capable; thus, the MPTCP session MUST operate as a regular, single-path TCP. If a SYN does not contain a MP_CAPABLE option, the SYN/ACK MUST NOT contain one in response. If the third packet (the ACK) does not contain the MP_CAPABLE option, then the session MUST fall back to operating as a regular, single-path TCP. This is to maintain compatibility with middleboxes on the path that drop some or all TCP options. Note that an implementation MAY choose to attempt sending MPTCP options more than one time before making this decision to operate as regular TCP (see Section 3.9).

If the SYN packets are unacknowledged, it is up to local policy to decide how to respond. It is expected that a sender will eventually fall back to single-path TCP (i.e., without the MP_CAPABLE option) in order to work around middleboxes that may drop packets with unknown options; however, the number of multipath-capable attempts that are made first will be up to local policy. It is possible that MPTCP and non-MPTCP SYNs could get reordered in the network. Therefore, the final state is inferred from the presence or absence of the MP_CAPABLE option in the third packet of the TCP handshake. If this option is not present, the connection SHOULD fall back to regular TCP, as documented in Section 3.7.

The initial data sequence number on an MPTCP connection is generated from the key. The algorithm for IDSN generation is also determined from the negotiated authentication algorithm. In this specification, with only the SHA-256 algorithm specified and selected, the IDSN of a host MUST be the least significant 64 bits of the SHA-256 hash of its key, i.e., $IDSN-A = Hash(Key-A)$ and $IDSN-B = Hash(Key-B)$. This deterministic generation of the IDSN allows a receiver to ensure that there are no gaps in sequence space at the start of the connection. The SYN with MP_CAPABLE occupies the first octet of data sequence space, although this does not need to be acknowledged at the connection level until the first data is sent (see Section 3.3).

3.2. Starting a New Subflow

Once an MPTCP connection has begun with the MP_CAPABLE exchange, further subflows can be added to the connection. Hosts have knowledge of their own address(es), and can become aware of the other host's addresses through signaling exchanges as described in Section 3.4. Using this knowledge, a host can initiate a new subflow over a currently unused pair of addresses. It is permitted for either host in a connection to initiate the creation of a new subflow, but it is expected that this will normally be the original connection initiator (see Section 3.9 for heuristics).

A new subflow is started as a normal TCP SYN/ACK exchange. The Join Connection (MP_JOIN) MPTCP option is used to identify the connection to be joined by the new subflow. It uses keying material that was exchanged in the initial MP_CAPABLE handshake (Section 3.1), and that handshake also negotiates the crypto algorithm in use for the MP_JOIN handshake.

This section specifies the behavior of MP_JOIN using the HMAC-SHA256 algorithm. An MP_JOIN option is present in the SYN, SYN/ACK, and ACK of the three-way handshake, although in each case with a different format.

In the first MP_JOIN on the SYN packet, illustrated in Figure 5, the initiator sends a token, random number, and address ID.

The token is used to identify the MPTCP connection and is a cryptographic hash of the receiver's key, as exchanged in the initial MP_CAPABLE handshake (Section 3.1). In this specification, the tokens presented in this option are generated by the SHA-256 [RFC6234] algorithm, truncated to the most significant 32 bits. The token included in the MP_JOIN option is the token that the receiver of the packet uses to identify this connection; i.e., Host A will send Token-B (which is generated from Key-B). Note that the hash generation algorithm can be overridden by the choice of cryptographic handshake algorithm, as defined in Section 3.1.

The MP_JOIN SYN sends not only the token (which is static for a connection) but also random numbers (nonces) that are used to prevent replay attacks on the authentication method. Recommendations for the generation of random numbers for this purpose are given in [RFC4086].

The MP_JOIN option includes an "Address ID". This is an identifier generated by the sender of the option, used to identify the source address of this packet, even if the IP header has been changed in transit by a middlebox. The numeric value of this field is generated by the sender and must map uniquely to a source IP address for the

sending host. The Address ID allows address removal (Section 3.4.2) without needing to know what the source address at the receiver is, thus allowing address removal through NATs. The Address ID also allows correlation between new subflow setup attempts and address signaling (Section 3.4.1), to prevent setting up duplicate subflows on the same path, if an MP_JOIN and ADD_ADDR are sent at the same time.

The Address IDs of the subflow used in the initial SYN exchange of the first subflow in the connection are implicit, and have the value zero. A host MUST store the mappings between Address IDs and addresses both for itself and the remote host. An implementation will also need to know which local and remote Address IDs are associated with which established subflows, for when addresses are removed from a local or remote host.

The MP_JOIN option on packets with the SYN flag set also includes 4 bits of flags, 3 of which are currently reserved and MUST be set to zero by the sender. The final bit, labeled "B", indicates whether the sender of this option wishes this subflow to be used as a backup path (B=1) in the event of failure of other paths, or whether it wants it to be used as part of the connection immediately. By setting B=1, the sender of the option is requesting the other host to only send data on this subflow if there are no available subflows where B=0. Subflow policy is discussed in more detail in Section 3.3.8.

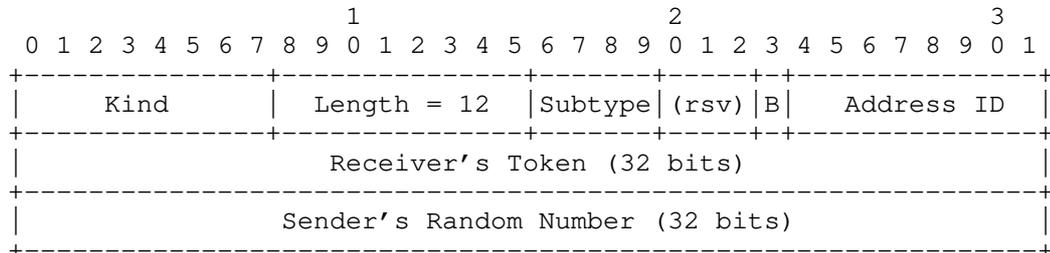


Figure 5: Join Connection (MP_JOIN) Option (for Initial SYN)

When receiving a SYN with an MP_JOIN option that contains a valid token for an existing MPTCP connection, the recipient SHOULD respond with a SYN/ACK also containing an MP_JOIN option containing a random number and a truncated (leftmost 64 bits) Hash-based Message Authentication Code (HMAC). This version of the option is shown in Figure 6. If the token is unknown, or the host wants to refuse subflow establishment (for example, due to a limit on the number of subflows it will permit), the receiver will send back a reset (RST) signal, analogous to an unknown port in TCP, containing a MP_TCP_RST

option (Section 3.6) with a "MPTCP specific error" reason code. Although calculating an HMAC requires cryptographic operations, it is believed that the 32-bit token in the MP_JOIN SYN gives sufficient protection against blind state exhaustion attacks; therefore, there is no need to provide mechanisms to allow a responder to operate statelessly at the MP_JOIN stage.

An HMAC is sent by both hosts -- by the initiator (Host A) in the third packet (the ACK) and by the responder (Host B) in the second packet (the SYN/ACK). Doing the HMAC exchange at this stage allows both hosts to have first exchanged random data (in the first two SYN packets) that is used as the "message". This specification defines that HMAC as defined in [RFC2104] is used, along with the SHA-256 hash algorithm [RFC6234], and that the output is truncated to the leftmost 160 bits (20 octets). Due to option space limitations, the HMAC included in the SYN/ACK is truncated to the leftmost 64 bits, but this is acceptable since random numbers are used; thus, an attacker only has one chance to correctly guess the HMAC that matches the random number previously sent by the peer (if the HMAC is incorrect, the TCP connection is closed, so a new MP_JOIN negotiation with a new random number is required).

The initiator's authentication information is sent in its first ACK (the third packet of the handshake), as shown in Figure 7. This data needs to be sent reliably, since it is the only time this HMAC is sent; therefore, receipt of this packet MUST trigger a regular TCP ACK in response, and the packet MUST be retransmitted if this ACK is not received. In other words, sending the ACK/MP_JOIN packet places the subflow in the PRE_ESTABLISHED state, and it moves to the ESTABLISHED state only on receipt of an ACK from the receiver. It is not permitted to send data while in the PRE_ESTABLISHED state. The reserved bits in this option MUST be set to zero by the sender.

The key for the HMAC algorithm, in the case of the message transmitted by Host A, will be Key-A followed by Key-B, and in the case of Host B, Key-B followed by Key-A. These are the keys that were exchanged in the original MP_CAPABLE handshake. The "message" for the HMAC algorithm in each case is the concatenations of random number for each host (denoted by R): for Host A, R-A followed by R-B; and for Host B, R-B followed by R-A.

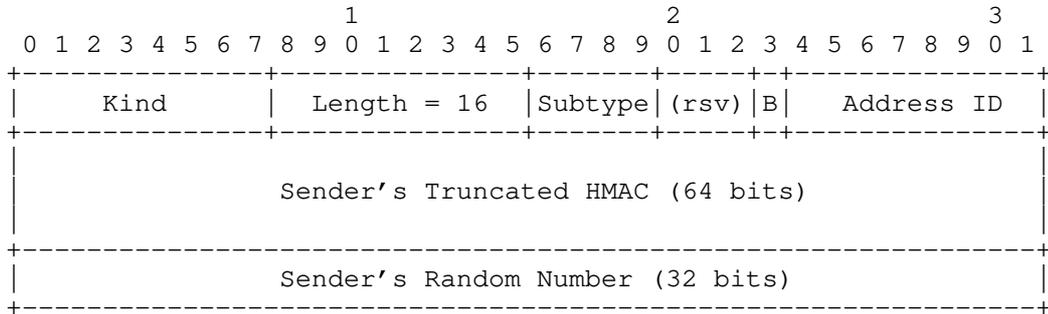


Figure 6: Join Connection (MP_JOIN) Option (for Responding SYN/ACK)

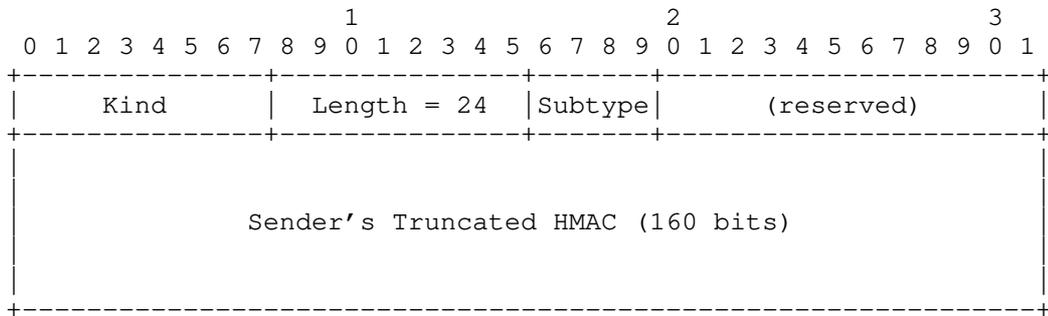
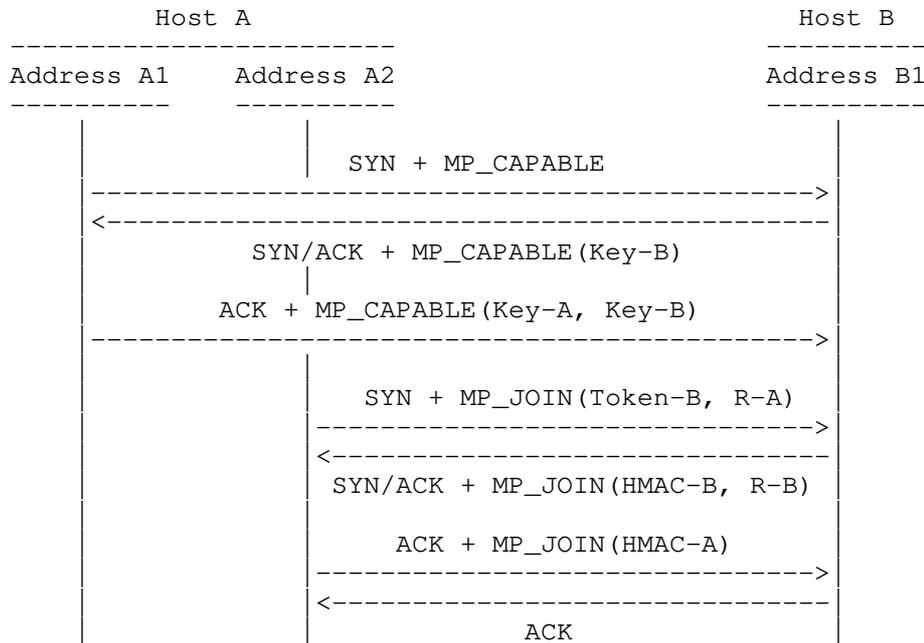


Figure 7: Join Connection (MP_JOIN) Option (for Third ACK)

These various MPTCP options fit together to enable authenticated subflow setup as illustrated in Figure 8.



HMAC-A = HMAC (Key=(Key-A+Key-B), Msg=(R-A+R-B))
 HMAC-B = HMAC (Key=(Key-B+Key-A), Msg=(R-B+R-A))

Figure 8: Example Use of MPTCP Authentication

If the token received at Host B is unknown or local policy prohibits the acceptance of the new subflow, the recipient MUST respond with a TCP RST for the subflow. If appropriate, a MP_TCPRST option with a "Administratively prohibited" reason code (Section 3.6) should be included.

If the token is accepted at Host B, but the HMAC returned to Host A does not match the one expected, Host A MUST close the subflow with a TCP RST. In this, and all following cases of sending a RST in this section, the sender SHOULD send a MP_TCPRST option (Section 3.6) on this RST packet with the reason code for a "MPTCP specific error".

If Host B does not receive the expected HMAC, or the MP_JOIN option is missing from the ACK, it MUST close the subflow with a TCP RST.

If the HMACs are verified as correct, then both hosts have verified each other as being the same peers as existed at the start of the connection, and they have agreed of which connection this subflow will become a part.

If the SYN/ACK as received at Host A does not have an MP_JOIN option, Host A MUST close the subflow with a TCP RST.

This covers all cases of the loss of an MP_JOIN. In more detail, if MP_JOIN is stripped from the SYN on the path from A to B, and Host B does not have a listener on the relevant port, it will respond with a RST in the normal way. If in response to a SYN with an MP_JOIN option, a SYN/ACK is received without the MP_JOIN option (either since it was stripped on the return path, or it was stripped on the outgoing path but Host B responded as if it were a new regular TCP session), then the subflow is unusable and Host A MUST close it with a RST.

Note that additional subflows can be created between any pair of ports (but see Section 3.9 for heuristics); no explicit application-level accept calls or bind calls are required to open additional subflows. To associate a new subflow with an existing connection, the token supplied in the subflow's SYN exchange is used for demultiplexing. This then binds the 5-tuple of the TCP subflow to the local token of the connection. A consequence is that it is possible to allow any port pairs to be used for a connection.

Demultiplexing subflow SYNs MUST be done using the token; this is unlike traditional TCP, where the destination port is used for demultiplexing SYN packets. Once a subflow is set up, demultiplexing packets is done using the 5-tuple, as in traditional TCP. The 5-tuples will be mapped to the local connection identifier (token). Note that Host A will know its local token for the subflow even though it is not sent on the wire -- only the responder's token is sent.

3.3. General MPTCP Operation

This section discusses operation of MPTCP for data transfer. At a high level, an MPTCP implementation will take one input data stream from an application, and split it into one or more subflows, with sufficient control information to allow it to be reassembled and delivered reliably and in order to the recipient application. The following subsections define this behavior in detail.

The data sequence mapping and the Data ACK are signaled in the Data Sequence Signal (DSS) option (Figure 9). Either or both can be signaled in one DSS, depending on the flags set. The data sequence mapping defines how the sequence space on the subflow maps to the connection level, and the Data ACK acknowledges receipt of data at the connection level. These functions are described in more detail in the following two subsections.

checksum is not present when its use has been negotiated, the receiver MUST close the subflow with a RST as it is considered broken. In both cases, this RST SHOULD be accompanied with a MP_TCPRST option (Section 3.6) with the reason code for a "MPTCP specific error".

3.3.1. Data Sequence Mapping

The data stream as a whole can be reassembled through the use of the data sequence mapping components of the DSS option (Figure 9), which define the mapping from the subflow sequence number to the data sequence number. This is used by the receiver to ensure in-order delivery to the application layer. Meanwhile, the subflow-level sequence numbers (i.e., the regular sequence numbers in the TCP header) have subflow-only relevance. It is expected (but not mandated) that SACK [RFC2018] is used at the subflow level to improve efficiency.

The data sequence mapping specifies a mapping from subflow sequence space to data sequence space. This is expressed in terms of starting sequence numbers for the subflow and the data level, and a length of bytes for which this mapping is valid. This explicit mapping for a range of data was chosen rather than per-packet signaling to assist with compatibility with situations where TCP/IP segmentation or coalescing is undertaken separately from the stack that is generating the data flow (e.g., through the use of TCP segmentation offloading on network interface cards, or by middleboxes such as performance enhancing proxies). It also allows a single mapping to cover many packets, which may be useful in bulk transfer situations.

A mapping is fixed, in that the subflow sequence number is bound to the data sequence number after the mapping has been processed. A sender MUST NOT change this mapping after it has been declared; however, the same data sequence number can be mapped to by different subflows for retransmission purposes (see Section 3.3.6). This would also permit the same data to be sent simultaneously on multiple subflows for resilience or efficiency purposes, especially in the case of lossy links. Although the detailed specification of such operation is outside the scope of this document, an implementation SHOULD treat the first data that is received at a subflow for the data sequence space as that which should be delivered to the application, and any later data for that sequence space SHOULD be ignored.

The data sequence number is specified as an absolute value, whereas the subflow sequence numbering is relative (the SYN at the start of the subflow has relative subflow sequence number 0). This is to allow middleboxes to change the initial sequence number of a subflow,

such as firewalls that undertake Initial Sequence Number (ISN) randomization.

The data sequence mapping also contains a checksum of the data that this mapping covers, if use of checksums has been negotiated at the MP_CAPABLE exchange. Checksums are used to detect if the payload has been adjusted in any way by a non-MPTCP-aware middlebox. If this checksum fails, it will trigger a failure of the subflow, or a fallback to regular TCP, as documented in Section 3.7, since MPTCP can no longer reliably know the subflow sequence space at the receiver to build data sequence mappings. Without checksumming enabled, corrupt data may be delivered to the application if a middlebox alters segment boundaries, alters content, or does not deliver all segments covered by a data sequence mapping. It is therefore RECOMMENDED to use checksumming unless it is known the network path contains no such devices.

The checksum algorithm used is the standard TCP checksum [RFC0793], operating over the data covered by this mapping, along with a pseudo-header as shown in Figure 10.

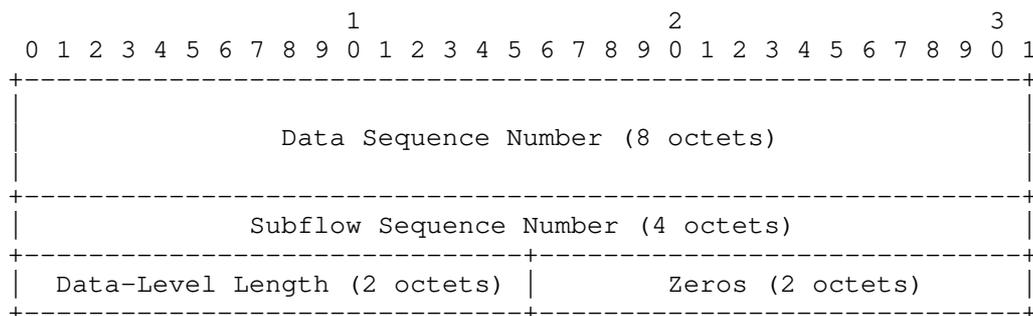


Figure 10: Pseudo-Header for DSS Checksum

Note that the data sequence number used in the pseudo-header is always the 64-bit value, irrespective of what length is used in the DSS option itself. The standard TCP checksum algorithm has been chosen since it will be calculated anyway for the TCP subflow, and if calculated first over the data before adding the pseudo-headers, it only needs to be calculated once. Furthermore, since the TCP checksum is additive, the checksum for a DSN_MAP can be constructed by simply adding together the checksums for the data of each constituent TCP segment, and adding the checksum for the DSS pseudo-header.

Note that checksumming relies on the TCP subflow containing contiguous data; therefore, a TCP subflow MUST NOT use the Urgent

Pointer to interrupt an existing mapping. Further note, however, that if Urgent data is received on a subflow, it SHOULD be mapped to the data sequence space and delivered to the application analogous to Urgent data in regular TCP.

To avoid possible deadlock scenarios, subflow-level processing should be undertaken separately from that at connection level. Therefore, even if a mapping does not exist from the subflow space to the data-level space, the data SHOULD still be ACKed at the subflow (if it is in-window). This data cannot, however, be acknowledged at the data level (Section 3.3.2) because its data sequence numbers are unknown. Implementations MAY hold onto such unmapped data for a short while in the expectation that a mapping will arrive shortly. Such unmapped data cannot be counted as being within the connection level receive window because this is relative to the data sequence numbers, so if the receiver runs out of memory to hold this data, it will have to be discarded. If a mapping for that subflow-level sequence space does not arrive within a receive window of data, that subflow SHOULD be treated as broken, closed with a RST, and any unmapped data silently discarded.

Data sequence numbers are always 64-bit quantities, and MUST be maintained as such in implementations. If a connection is progressing at a slow rate, so protection against wrapped sequence numbers is not required, then an implementation MAY include just the lower 32 bits of the data sequence number in the data sequence mapping and/or Data ACK as an optimization, and an implementation can make this choice independently for each packet. An implementation MUST be able to receive and process both 64-bit or 32-bit sequence number values, but it is not required that an implementation is able to send both.

An implementation MUST send the full 64-bit data sequence number if it is transmitting at a sufficiently high rate that the 32-bit value could wrap within the Maximum Segment Lifetime (MSL) [RFC7323]. The lengths of the DSNs used in these values (which may be different) are declared with flags in the DSS option. Implementations MUST accept a 32-bit DSN and implicitly promote it to a 64-bit quantity by incrementing the upper 32 bits of sequence number each time the lower 32 bits wrap. A sanity check MUST be implemented to ensure that a wrap occurs at an expected time (e.g., the sequence number jumps from a very high number to a very low number) and is not triggered by out-of-order packets.

As with the standard TCP sequence number, the data sequence number should not start at zero, but at a random value to make blind session hijacking harder. This specification requires setting the initial data sequence number (IDSN) of each host to the least significant 64

bits of the SHA-256 hash of the host's key, as described in Section 3.1. This is required also in order for the receiver to know what the expected IDSN is, and thus determine if any initial connection-level packets are missing; this is particularly relevant if two subflows start transmitting simultaneously.

A data sequence mapping does not need to be included in every MPTCP packet, as long as the subflow sequence space in that packet is covered by a mapping known at the receiver. This can be used to reduce overhead in cases where the mapping is known in advance; one such case is when there is a single subflow between the hosts, another is when segments of data are scheduled in larger than packet-sized chunks.

An "infinite" mapping can be used to fall back to regular TCP by mapping the subflow-level data to the connection-level data for the remainder of the connection (see Section 3.7). This is achieved by setting the Data-Level Length field of the DSS option to the reserved value of 0. The checksum, in such a case, will also be set to zero.

3.3.2. Data Acknowledgments

To provide full end-to-end resilience, MPTCP provides a connection-level acknowledgment, to act as a cumulative ACK for the connection as a whole. This is the "Data ACK" field of the DSS option (Figure 9). The Data ACK is analogous to the behavior of the standard TCP cumulative ACK -- indicating how much data has been successfully received (with no holes). This is in comparison to the subflow-level ACK, which acts analogous to TCP SACK, given that there may still be holes in the data stream at the connection level. The Data ACK specifies the next data sequence number it expects to receive.

The Data ACK, as for the DSN, can be sent as the full 64-bit value, or as the lower 32 bits. If data is received with a 64-bit DSN, it MUST be acknowledged with a 64-bit Data ACK. If the DSN received is 32 bits, an implementation can choose whether to send a 32-bit or 64-bit Data ACK, and an implementation MUST accept either in this situation.

The Data ACK proves that the data, and all required MPTCP signaling, has been received and accepted by the remote end. One key use of the Data ACK signal is that it is used to indicate the left edge of the advertised receive window. As explained in Section 3.3.4, the receive window is shared by all subflows and is relative to the Data ACK. Because of this, an implementation MUST NOT use the RCV.WND field of a TCP segment at the connection level if it does not also carry a DSS option with a Data ACK field. Furthermore, separating

the connection-level acknowledgments from the subflow level allows processing to be done separately, and a receiver has the freedom to drop segments after acknowledgment at the subflow level, for example, due to memory constraints when many segments arrive out of order.

An MPTCP sender MUST NOT free data from the send buffer until it has been acknowledged by both a Data ACK received on any subflow and at the subflow level by all subflows on which the data was sent. The former condition ensures liveness of the connection and the latter condition ensures liveness and self-consistence of a subflow when data needs to be retransmitted. Note, however, that if some data needs to be retransmitted multiple times over a subflow, there is a risk of blocking the sending window. In this case, the MPTCP sender can decide to terminate the subflow that is behaving badly by sending a RST, using an appropriate MP_TCP_RST (Section 3.6) error code.

The Data ACK MAY be included in all segments; however, optimizations SHOULD be considered in more advanced implementations, where the Data ACK is present in segments only when the Data ACK value advances, and this behavior MUST be treated as valid. This behavior ensures the sender buffer is freed, while reducing overhead when the data transfer is unidirectional.

3.3.3. Closing a Connection

In regular TCP, a FIN announces the receiver that the sender has no more data to send. In order to allow subflows to operate independently and to keep the appearance of TCP over the wire, a FIN in MPTCP only affects the subflow on which it is sent. This allows nodes to exercise considerable freedom over which paths are in use at any one time. The semantics of a FIN remain as for regular TCP; i.e., it is not until both sides have ACKed each other's FINs that the subflow is fully closed.

When an application calls close() on a socket, this indicates that it has no more data to send; for regular TCP, this would result in a FIN on the connection. For MPTCP, an equivalent mechanism is needed, and this is referred to as the DATA_FIN.

A DATA_FIN is an indication that the sender has no more data to send, and as such can be used to verify that all data has been successfully received. A DATA_FIN, as with the FIN on a regular TCP connection, is a unidirectional signal.

The DATA_FIN is signaled by setting the 'F' flag in the Data Sequence Signal option (Figure 9) to 1. A DATA_FIN occupies 1 octet (the final octet) of the connection-level sequence space. Note that the DATA_FIN is included in the Data-Level Length, but not at the subflow

level: for example, a segment with DSN 80, and Data-Level Length 11, with DATA_FIN set, would map 10 octets from the subflow into data sequence space 80-89, the DATA_FIN is DSN 90; therefore, this segment including DATA_FIN would be acknowledged with a DATA_ACK of 91.

Note that when the DATA_FIN is not attached to a TCP segment containing data, the Data Sequence Signal MUST have a subflow sequence number of 0, a Data-Level Length of 1, and the data sequence number that corresponds with the DATA_FIN itself. The checksum in this case will only cover the pseudo-header.

A DATA_FIN has the semantics and behavior as a regular TCP FIN, but at the connection level. Notably, it is only DATA_ACKed once all data has been successfully received at the connection level. Note, therefore, that a DATA_FIN is decoupled from a subflow FIN. It is only permissible to combine these signals on one subflow if there is no data outstanding on other subflows. Otherwise, it may be necessary to retransmit data on different subflows. Essentially, a host MUST NOT close all functioning subflows unless it is safe to do so, i.e., until all outstanding data has been DATA_ACKed, or until the segment with the DATA_FIN flag set is the only outstanding segment.

Once a DATA_FIN has been acknowledged, all remaining subflows MUST be closed with standard FIN exchanges. Both hosts SHOULD send FINs on all subflows, as a courtesy to allow middleboxes to clean up state even if an individual subflow has failed. It is also encouraged to reduce the timeouts (Maximum Segment Lifetime) on subflows at end hosts after receiving a DATA_FIN. In particular, any subflows where there is still outstanding data queued (which has been retransmitted on other subflows in order to get the DATA_FIN acknowledged) MAY be closed with a RST with MP_TCP_RST (Section 3.6) error code for "too much outstanding data".

A connection is considered closed once both hosts' DATA_FINs have been acknowledged by DATA_ACKs.

As specified above, a standard TCP FIN on an individual subflow only shuts down the subflow on which it was sent. If all subflows have been closed with a FIN exchange, but no DATA_FIN has been received and acknowledged, the MPTCP connection is treated as closed only after a timeout. This implies that an implementation will have TIME_WAIT states at both the subflow and connection levels (see Appendix D). This permits "break-before-make" scenarios where connectivity is lost on all subflows before a new one can be re-established.

3.3.4. Receiver Considerations

Regular TCP advertises a receive window in each packet, telling the sender how much data the receiver is willing to accept past the cumulative ack. The receive window is used to implement flow control, throttling down fast senders when receivers cannot keep up.

MPTCP also uses a unique receive window, shared between the subflows. The idea is to allow any subflow to send data as long as the receiver is willing to accept it. The alternative, maintaining per subflow receive windows, could end up stalling some subflows while others would not use up their window.

The receive window is relative to the `DATA_ACK`. As in TCP, a receiver **MUST NOT** shrink the right edge of the receive window (i.e., `DATA_ACK + receive window`). The receiver will use the data sequence number to tell if a packet should be accepted at the connection level.

When deciding to accept packets at subflow level, regular TCP checks the sequence number in the packet against the allowed receive window. With multipath, such a check is done using only the connection-level window. A sanity check **SHOULD** be performed at subflow level to ensure that the subflow and mapped sequence numbers meet the following test: $SSN - SUBFLOW_ACK \leq DSN - DATA_ACK$, where `SSN` is the subflow sequence number of the received packet and `SUBFLOW_ACK` is the `RCV.NXT` (next expected sequence number) of the subflow (with the equivalent connection-level definitions for `DSN` and `DATA_ACK`).

In regular TCP, once a segment is deemed in-window, it is put either in the in-order receive queue or in the out-of-order queue. In Multipath TCP, the same happens but at the connection level: a segment is placed in the connection level in-order or out-of-order queue if it is in-window at both connection and subflow levels. The stack still has to remember, for each subflow, which segments were received successfully so that it can ACK them at subflow level appropriately. Typically, this will be implemented by keeping per subflow out-of-order queues (containing only message headers, not the payloads) and remembering the value of the cumulative ACK.

It is important for implementers to understand how large a receiver buffer is appropriate. The lower bound for full network utilization is the maximum bandwidth-delay product of any one of the paths. However, this might be insufficient when a packet is lost on a slower subflow and needs to be retransmitted (see Section 3.3.6). A tight upper bound would be the maximum round-trip time (RTT) of any path multiplied by the total bandwidth available across all paths. This permits all subflows to continue at full speed while a packet is

fast-retransmitted on the maximum RTT path. Even this might be insufficient to maintain full performance in the event of a retransmit timeout on the maximum RTT path. It is for future study to determine the relationship between retransmission strategies and receive buffer sizing.

3.3.5. Sender Considerations

The sender remembers receiver window advertisements from the receiver. It should only update its local receive window values when the largest sequence number allowed (i.e., `DATA_ACK` + receive window) increases, on the receipt of a `DATA_ACK`. This is important to allow using paths with different RTTs, and thus different feedback loops.

MPTCP uses a single receive window across all subflows, and if the receive window was guaranteed to be unchanged end-to-end, a host could always read the most recent receive window value. However, some classes of middleboxes may alter the TCP-level receive window. Typically, these will shrink the offered window, although for short periods of time it may be possible for the window to be larger (however, note that this would not continue for long periods since ultimately the middlebox must keep up with delivering data to the receiver). Therefore, if receive window sizes differ on multiple subflows, when sending data MPTCP SHOULD take the largest of the most recent window sizes as the one to use in calculations. This rule is implicit in the requirement not to reduce the right edge of the window.

The sender MUST also remember the receive windows advertised by each subflow. The allowed window for subflow *i* is (`ack_i`, `ack_i` + `rcv_wnd_i`), where `ack_i` is the subflow-level cumulative ACK of subflow *i*. This ensures data will not be sent to a middlebox unless there is enough buffering for the data.

Putting the two rules together, we get the following: a sender is allowed to send data segments with data-level sequence numbers between (`DATA_ACK`, `DATA_ACK` + `receive_window`). Each of these segments will be mapped onto subflows, as long as subflow sequence numbers are in the allowed windows for those subflows. Note that subflow sequence numbers do not generally affect flow control if the same receive window is advertised across all subflows. They will perform flow control for those subflows with a smaller advertised receive window.

The send buffer MUST, at a minimum, be as big as the receive buffer, to enable the sender to reach maximum throughput.

3.3.6. Reliability and Retransmissions

The data sequence mapping allows senders to resend data with the same data sequence number on a different subflow. When doing this, a host **MUST** still retransmit the original data on the original subflow, in order to preserve the subflow integrity (middleboxes could replay old data, and/or could reject holes in subflows), and a receiver will ignore these retransmissions. While this is clearly suboptimal, for compatibility reasons this is sensible behavior. Optimizations could be negotiated in future versions of this protocol. Note also that this property would also permit a sender to always send the same data, with the same data sequence number, on multiple subflows, if desired for reliability reasons.

This protocol specification does not mandate any mechanisms for handling retransmissions, and much will be dependent upon local policy (as discussed in Section 3.3.8). One can imagine aggressive connection-level retransmissions policies where every packet lost at subflow level is retransmitted on a different subflow (hence, wasting bandwidth but possibly reducing application-to-application delays), or conservative retransmission policies where connection-level retransmits are only used after a few subflow-level retransmission timeouts occur.

It is envisaged that a standard connection-level retransmission mechanism would be implemented around a connection-level data queue: all segments that haven't been `DATA_ACK`ed are stored. A timer is set when the head of the connection-level is `ACK`ed at subflow level but its corresponding data is not `ACK`ed at data level. This timer will guard against failures in retransmission by middleboxes that proactively `ACK` data.

The sender **MUST** keep data in its send buffer as long as the data has not been acknowledged at both connection level and on all subflows on which it has been sent. In this way, the sender can always retransmit the data if needed, on the same subflow or on a different one. A special case is when a subflow fails: the sender will typically resend the data on other working subflows after a timeout, and will keep trying to retransmit the data on the failed subflow too. The sender will declare the subflow failed after a predefined upper bound on retransmissions is reached (which **MAY** be lower than the usual TCP limits of the Maximum Segment Life), or on the receipt of an ICMP error, and only then delete the outstanding data segments.

If multiple retransmissions are triggered that indicate that a subflow performs badly, this **MAY** lead to a host resetting the subflow with a RST. However, additional research is required to understand the heuristics of how and when to reset underperforming subflows.

For example, a highly asymmetric path may be misdiagnosed as underperforming. A RST for this purpose SHOULD be accompanied with an "Unacceptable performance" MP_TCPRST option (Section 3.6).

3.3.7. Congestion Control Considerations

Different subflows in an MPTCP connection have different congestion windows. To achieve fairness at bottlenecks and resource pooling, it is necessary to couple the congestion windows in use on each subflow, in order to push most traffic to uncongested links. One algorithm for achieving this is presented in [RFC6356]; the algorithm does not achieve perfect resource pooling but is "safe" in that it is readily deployable in the current Internet. By this, we mean that it does not take up more capacity on any one path than if it was a single path flow using only that route, so this ensures fair coexistence with single-path TCP at shared bottlenecks.

It is foreseeable that different congestion controllers will be implemented for MPTCP, each aiming to achieve different properties in the resource pooling/fairness/stability design space, as well as those for achieving different properties in quality of service, reliability, and resilience.

Regardless of the algorithm used, the design of the MPTCP protocol aims to provide the congestion control implementations sufficient information to take the right decisions; this information includes, for each subflow, which packets were lost and when.

3.3.8. Subflow Policy

Within a local MPTCP implementation, a host may use any local policy it wishes to decide how to share the traffic to be sent over the available paths.

In the typical use case, where the goal is to maximize throughput, all available paths will be used simultaneously for data transfer, using coupled congestion control as described in [RFC6356]. It is expected, however, that other use cases will appear.

For instance, a possibility is an 'all-or-nothing' approach, i.e., have a second path ready for use in the event of failure of the first path, but alternatives could include entirely saturating one path before using an additional path (the 'overflow' case). Such choices would be most likely based on the monetary cost of links, but may also be based on properties such as the delay or jitter of links, where stability (of delay or bandwidth) is more important than throughput. Application requirements such as these are discussed in detail in [RFC6897].

The ability to make effective choices at the sender requires full knowledge of the path "cost", which is unlikely to be the case. It would be desirable for a receiver to be able to signal their own preferences for paths, since they will often be the multihomed party, and may have to pay for metered incoming bandwidth.

To enable this, the MP_JOIN option (see Section 3.2) contains the 'B' bit, which allows a host to indicate to its peer that this path should be treated as a backup path to use only in the event of failure of other working subflows (i.e., a subflow where the receiver has indicated B=1 SHOULD NOT be used to send data unless there are no usable subflows where B=0).

In the event that the available set of paths changes, a host may wish to signal a change in priority of subflows to the peer (e.g., a subflow that was previously set as backup should now take priority over all remaining subflows). Therefore, the MP_PRIO option, shown in Figure 11, can be used to change the 'B' flag of the subflow on which it is sent.

Another use of the MP_PRIO option is to set the 'B' flag on a subflow to cleanly retire its use before closing it and removing it with REMOVE_ADDR Section 3.4.2, for example to support make-before-break session continuity, where new subflows are added before the previously used ones are closed.

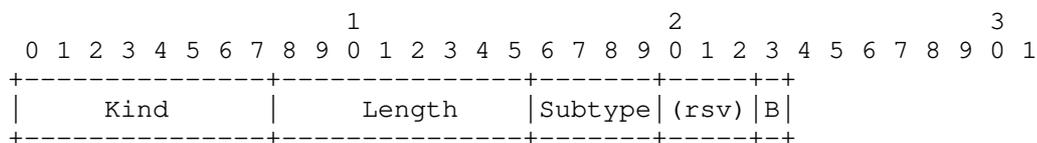


Figure 11: Change Subflow Priority (MP_PRIO) Option

It should be noted that the backup flag is a request from a data receiver to a data sender only, and the data sender SHOULD adhere to these requests. A host cannot assume that the data sender will do so, however, since local policies -- or technical difficulties -- may override MP_PRIO requests. Note also that this signal applies to a single direction, and so the sender of this option could choose to continue using the subflow to send data even if it has signaled B=1 to the other host.

3.4. Address Knowledge Exchange (Path Management)

We use the term "path management" to refer to the exchange of information about additional paths between hosts, which in this design is managed by multiple addresses at hosts. For more detail of

the architectural thinking behind this design, see the MPTCP Architecture document [RFC6182].

This design makes use of two methods of sharing such information, and both can be used on a connection. The first is the direct setup of new subflows, already described in Section 3.2, where the initiator has an additional address. The second method, described in the following subsections, signals addresses explicitly to the other host to allow it to initiate new subflows. The two mechanisms are complementary: the first is implicit and simple, while the explicit is more complex but is more robust. Together, the mechanisms allow addresses to change in flight (and thus support operation through NATs, since the source address need not be known), and also allow the signaling of previously unknown addresses, and of addresses belonging to other address families (e.g., both IPv4 and IPv6).

Here is an example of typical operation of the protocol:

- o An MPTCP connection is initially set up between address/port A1 of Host A and address/port B1 of Host B. If Host A is multihomed and multiaddressed, it can start an additional subflow from its address A2 to B1, by sending a SYN with a Join option from A2 to B1, using B's previously declared token for this connection. Alternatively, if B is multihomed, it can try to set up a new subflow from B2 to A1, using A's previously declared token. In either case, the SYN will be sent to the port already in use for the original subflow on the receiving host.
- o Simultaneously (or after a timeout), an ADD_ADDR option (Section 3.4.1) is sent on an existing subflow, informing the receiver of the sender's alternative address(es). The recipient can use this information to open a new subflow to the sender's additional address. In our example, A will send ADD_ADDR option informing B of address/port A2. The mix of using the SYN-based option and the ADD_ADDR option, including timeouts, is implementation specific and can be tailored to agree with local policy.
- o If subflow A2-B1 is successfully set up, Host B can use the Address ID in the Join option to correlate this with the ADD_ADDR option that will also arrive on an existing subflow; now B knows not to open A2-B1, ignoring the ADD_ADDR. Otherwise, if B has not received the A2-B1 MP_JOIN SYN but received the ADD_ADDR, it can try to initiate a new subflow from one or more of its addresses to address A2. This permits new sessions to be opened if one host is behind a NAT.

Other ways of using the two signaling mechanisms are possible; for instance, signaling addresses in other address families can only be done explicitly using the Add Address option.

3.4.1. Address Advertisement

The Add Address (ADD_ADDR) MPTCP option announces additional addresses (and optionally, ports) on which a host can be reached (Figure 12). This option can be used at any time during a connection, depending on when the sender wishes to enable multiple paths and/or when paths become available. As with all MPTCP signals, the receiver MUST undertake standard TCP validity checks, e.g. [RFC5961], before acting upon it.

Every address has an Address ID that can be used for uniquely identifying the address within a connection for address removal. The Address ID is also used to identify MP_JOIN options (see Section 3.2) relating to the same address, even when address translators are in use. The Address ID MUST uniquely identify the address for the sender of the option (within the scope of the connection), but the mechanism for allocating such IDs is implementation specific.

All address IDs learned via either MP_JOIN or ADD_ADDR SHOULD be stored by the receiver in a data structure that gathers all the Address ID to address mappings for a connection (identified by a token pair). In this way, there is a stored mapping between Address ID, observed source address, and token pair for future processing of control information for a connection. Note that an implementation MAY discard incoming address advertisements at will, for example, for avoiding updating mapping state, or because advertised addresses are of no use to it (for example, IPv6 addresses when it has IPv4 only). Therefore, a host MUST treat address advertisements as soft state, and it MAY choose to refresh advertisements periodically. Note also that an implementation MAY choose to cache these address advertisements even if they are not currently relevant but may be relevant in the future, such as IPv4 addresses when IPv6 connectivity is available but IPv4 is awaiting DHCP.

This option is shown in Figure 12. The illustration is sized for IPv4 addresses. For IPv6, the length of the address will be 16 octets (instead of 4).

The 2 octets that specify the TCP port number to use are optional and their presence can be inferred from the length of the option. Although it is expected that the majority of use cases will use the same port pairs as used for the initial subflow (e.g., port 80 remains port 80 on all subflows, as does the ephemeral port at the client), there may be cases (such as port-based load balancing) where

the explicit specification of a different port is required. If no port is specified, MPTCP SHOULD attempt to connect to the specified address on the same port as is already in use by the subflow on which the ADD_ADDR signal was sent; this is discussed in more detail in Section 3.9.

The Truncated HMAC present in this Option is the rightmost 64 bits of an HMAC, negotiated and calculated in the same way as for MP_JOIN as described in Section 3.2. For this specification of MPTCP, as there is only one hash algorithm option specified, this will be HMAC as defined in [RFC2104], using the SHA-256 hash algorithm [RFC6234]. In the same way as for MP_JOIN, the key for the HMAC algorithm, in the case of the message transmitted by Host A, will be Key-A followed by Key-B, and in the case of Host B, Key-B followed by Key-A. These are the keys that were exchanged in the original MP_CAPABLE handshake. The message for the HMAC is the Address ID, IP Address, and Port which precede the HMAC in the ADD_ADDR option. If the port is not present in the ADD_ADDR option, the HMAC message will nevertheless include two octets of value zero. The rationale for the HMAC is to prevent unauthorized entities from injecting ADD_ADDR signals in an attempt to hijack a connection. Note that additionally the presence of this HMAC prevents the address being changed in flight unless the key is known by an intermediary. If a host receives an ADD_ADDR option for which it cannot validate the HMAC, it SHOULD silently ignore the option.

A set of four flags are present after the subtype and before the Address ID. Only the rightmost bit - labelled 'E' - is assigned in this specification. The other bits are currently unassigned and MUST be set to zero by a sender and MUST be ignored by the receiver.

The 'E' flag exists to provide reliability for this option. Because this option will often be sent on pure ACKs, there is no guarantee of reliability. Therefore, a receiver receiving a fresh ADD_ADDR option (where E=0), will send the same option back to the sender, but not including the HMAC, and with E=1, to indicate receipt. The lack of this echo can be used by the initial ADD_ADDR sender to retransmit the ADD_ADDR according to local policy.

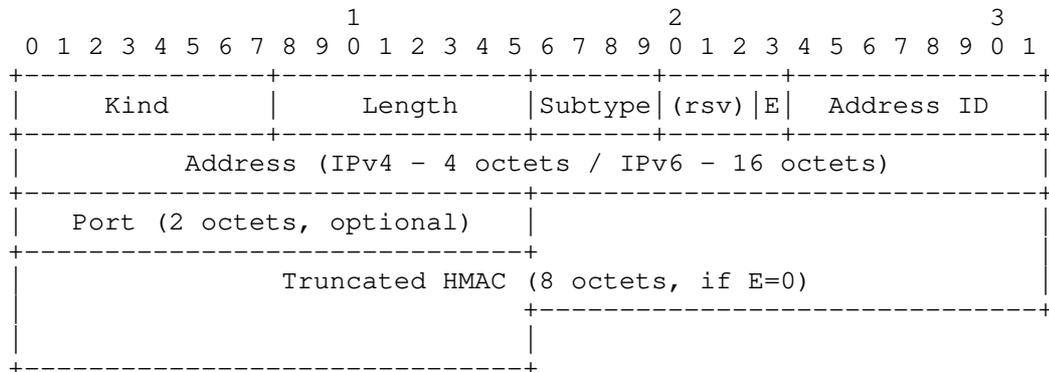


Figure 12: Add Address (ADD_ADDR) Option

Due to the proliferation of NATs, it is reasonably likely that one host may attempt to advertise private addresses [RFC1918]. It is not desirable to prohibit this, since there may be cases where both hosts have additional interfaces on the same private network, and a host MAY advertise such addresses. The MP_JOIN handshake to create a new subflow (Section 3.2) provides mechanisms to minimize security risks. The MP_JOIN message contains a 32-bit token that uniquely identifies the connection to the receiving host. If the token is unknown, the host will return with a RST. In the unlikely event that the token is valid at the receiving host, subflow setup will continue, but the HMAC exchange must occur for authentication. This will fail, and will provide sufficient protection against two unconnected hosts accidentally setting up a new subflow upon the signal of a private address. Further security considerations around the issue of ADD_ADDR messages that accidentally misdirect, or maliciously direct, new MP_JOIN attempts are discussed in Section 5.

A host that receives an ADD_ADDR but finds a connection set up to that IP address and port number is unsuccessful SHOULD NOT perform further connection attempts to this address/port combination for this connection. A sender that wants to trigger a new incoming connection attempt on a previously advertised address/port combination can therefore refresh ADD_ADDR information by sending the option again.

A host can therefore send an ADD_ADDR message with an already assigned Address ID, but the Address MUST be the same as previously assigned to this Address ID. A new ADD_ADDR may have the same, or different, port number. If the port number is different, the receiving host SHOULD try to set up a new subflow to this new address/port combination.

A host wishing to replace an existing Address ID MUST first remove the existing one (Section 3.4.2).

During normal MPTCP operation, it is unlikely that there will be sufficient TCP option space for ADD_ADDR to be included along with those for data sequence numbering (Section 3.3.1). Therefore, it is expected that an MPTCP implementation will send the ADD_ADDR option on separate ACKs. As discussed earlier, however, an MPTCP implementation MUST NOT treat duplicate ACKs with any MPTCP option, with the exception of the DSS option, as indications of congestion [RFC5681], and an MPTCP implementation SHOULD NOT send more than two duplicate ACKs in a row for signaling purposes.

3.4.2. Remove Address

If, during the lifetime of an MPTCP connection, a previously announced address becomes invalid (e.g., if the interface disappears, or an IPv6 address is no longer preferred), the affected host SHOULD announce this so that the peer can remove subflows related to this address. Even if an address is not in use by a MPTCP connection, if it has been previously announced, an implementation SHOULD announce its removal. A host MAY also choose to announce that a valid IP address should not be used any longer, for example for make-before-break session continuity.

This is achieved through the Remove Address (REMOVE_ADDR) option (Figure 13), which will remove a previously added address (or list of addresses) from a connection and terminate any subflows currently using that address.

For security purposes, if a host receives a REMOVE_ADDR option, it must ensure the affected path(s) are no longer in use before it instigates closure. The receipt of REMOVE_ADDR SHOULD first trigger the sending of a TCP keepalive [RFC1122] on the path, and if a response is received the path SHOULD NOT be removed. If the path is found to still be alive, the receiving host SHOULD no longer use the specified address for future connections, but it is the responsibility of the host which sent the REMOVE_ADDR to shut down the subflow. The requesting host MAY also use MP_PRIO (Section 3.3.8) to request a path is no longer used, before removal. Typical TCP validity tests on the subflow (e.g., ensuring sequence and ACK numbers are correct) MUST also be undertaken. An implementation can use indications of these test failures as part of intrusion detection or error logging.

The sending and receipt (if no keepalive response was received) of this message SHOULD trigger the sending of RSTs by both hosts on the

affected subflow(s) (if possible), as a courtesy to cleaning up middlebox state, before cleaning up any local state.

Address removal is undertaken by ID, so as to permit the use of NATs and other middleboxes that rewrite source addresses. If there is no address at the requested ID, the receiver will silently ignore the request.

A subflow that is still functioning MUST be closed with a FIN exchange as in regular TCP, rather than using this option. For more information, see Section 3.3.3.

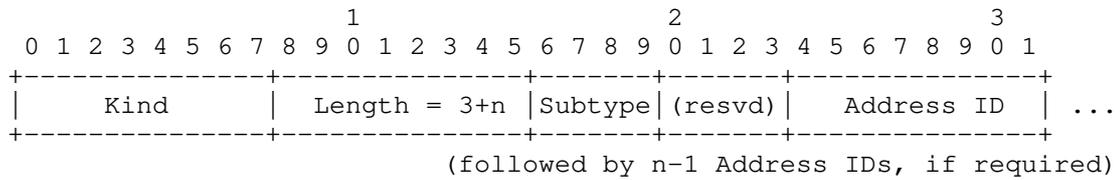


Figure 13: Remove Address (REMOVE_ADDR) Option

3.5. Fast Close

Regular TCP has the means of sending a reset (RST) signal to abruptly close a connection. With MPTCP, a regular RST only has the scope of the subflow and will only close the concerned subflow but not affect the remaining subflows. MPTCP's connection will stay alive at the data level, in order to permit break-before-make handover between subflows. It is therefore necessary to provide an MPTCP-level "reset" to allow the abrupt closure of the whole MPTCP connection, and this is the MP_FASTCLOSE option.

MP_FASTCLOSE is used to indicate to the peer that the connection will be abruptly closed and no data will be accepted anymore. The reasons for triggering an MP_FASTCLOSE are implementation specific. Regular TCP does not allow sending a RST while the connection is in a synchronized state [RFC0793]. Nevertheless, implementations allow the sending of a RST in this state, if, for example, the operating system is running out of resources. In these cases, MPTCP should send the MP_FASTCLOSE. This option is illustrated in Figure 14.

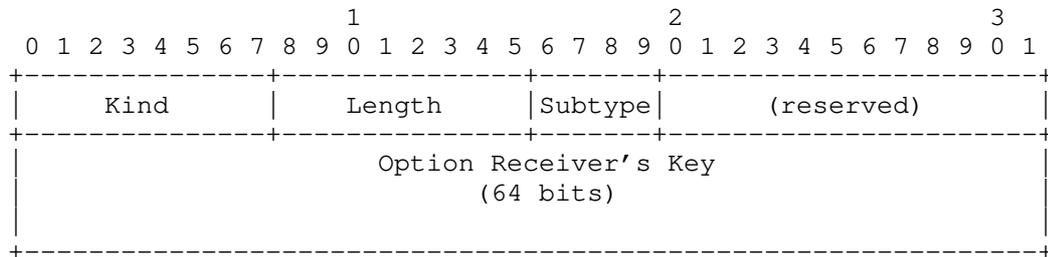


Figure 14: Fast Close (MP_FASTCLOSE) Option

If Host A wants to force the closure of an MPTCP connection, it has two different options:

- o Option A (ACK) : Host A sends an ACK containing the MP_FASTCLOSE option on one subflow, containing the key of Host B as declared in the initial connection handshake. On all the other subflows, Host A sends a regular TCP RST to close these subflows, and tears them down. Host A now enters FASTCLOSE_WAIT state.
- o Option R (RST) : Host A sends a RST containing the MP_FASTCLOSE option on all subflows, containing the key of Host B as declared in the initial connection handshake. Host A can tear the subflows and the connection down immediately.

If host A decides to force the closure by using Option A and sending an ACK with the MP_FASTCLOSE option, the connection shall proceed as follows:

- o Upon receipt of an ACK with MP_FASTCLOSE by Host B, containing the valid key, Host B answers on the same subflow with a TCP RST and tears down all subflows also through sending TCP RST signals. Host B can now close the whole MPTCP connection (it transitions directly to CLOSED state).
- o As soon as Host A has received the TCP RST on the remaining subflow, it can close this subflow and tear down the whole connection (transition from FASTCLOSE_WAIT to CLOSED states). If Host A receives an MP_FASTCLOSE instead of a TCP RST, both hosts attempted fast closure simultaneously. Host A should reply with a TCP RST and tear down the connection.
- o If Host A does not receive a TCP RST in reply to its MP_FASTCLOSE after one retransmission timeout (RTO) (the RTO of the subflow where the MP_FASTCLOSE has been sent), it SHOULD retransmit the MP_FASTCLOSE. The number of retransmissions SHOULD be limited to avoid this connection from being retained for a long time, but

this limit is implementation specific. A RECOMMENDED number is 3. If no TCP RST is received in response, Host A SHOULD send a TCP RST with the MP_FASTCLOSE option itself when it releases state in order to clear any remaining state at middleboxes.

If however host A decides to force the closure by using Option R and sending a RST with the MP_FASTCLOSE option, Host B will act as follows: Upon receipt of a RST with MP_FASTCLOSE, containing the valid key, Host B tears down all subflows by sending a TCP RST. Host B can now close the whole MPTCP connection (it transitions directly to CLOSED state).

3.6. Subflow Reset

An implementation of MPTCP may also need to send a regular TCP RST to force the closure of a subflow. A host sends a TCP RST in order to close a subflow or reject an attempt to open a subflow (MP_JOIN). In order to inform the receiving host why a subflow is being closed or rejected, the TCP RST packet MAY include the MP_TCPRST Option. The host MAY use this information to decide, for example, whether it tries to re-establish the subflow immediately, later, or never.

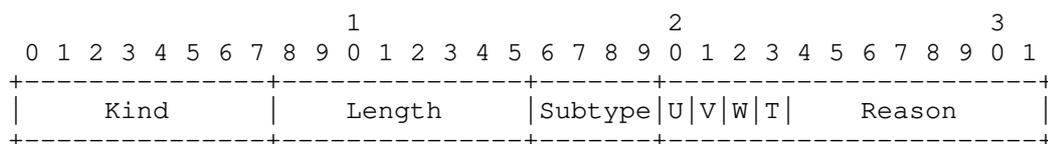


Figure 15: TCP RST Reason (MP_TCPRST) Option

The MP_TCPRST option contains a reason code that allows the sender of the option to provide more information about the reason for the termination of the subflow. Using 12 bits of option space, the first four bits are reserved for flags (only one of which is currently defined), and the remaining octet is used to express a reason code for this subflow termination, from which a receiver MAY infer information about the usability of this path.

The "T" flag is used by the sender to indicate whether the error condition that is reported is Transient (T bit set to 1) or Permanent (T bit set to 0). If the error condition is considered to be Transient by the sender of the RST segment, the recipient of this segment MAY try to reestablish a subflow for this connection over the failed path. The time at which a receiver may try to re-establish this is implementation-specific, but SHOULD take into account the properties of the failure defined by the following reason code. If the error condition is considered to be permanent, the receiver of the RST segment SHOULD NOT try to reestablish a subflow for this

connection over this path. The "U", "V" and "W" flags are not defined by this specification and are reserved for future use. An implementation of this specification MUST set these flags to 0, and a receiver MUST ignore them.

The "Reason" code is an 8-bit field that indicates the reason for the termination of the subflow. The following codes are defined in this document:

- o Unspecified error (code 0x0). This is the default error implying the subflow is no longer available. The presence of this option shows that the RST was generated by a MPTCP-aware device.
- o MPTCP specific error (code 0x01). An error has been detected in the processing of MPTCP options. This is the usual reason code to return in the cases where a RST is being sent to close a subflow for reasons of an invalid response.
- o Lack of resources (code 0x02). This code indicates that the sending host does not have enough resources to support the terminated subflow.
- o Administratively prohibited (code 0x03). This code indicates that the requested subflow is prohibited by the policies of the sending host.
- o Too much outstanding data (code 0x04). This code indicates that there is an excessive amount of data that need to be transmitted over the terminated subflow while having already been acknowledged over one or more other subflows. This may occur if a path has been unavailable for a short period and it is more efficient to reset and start again than it is to retransmit the queued data.
- o Unacceptable performance (code 0x05). This code indicates that the performance of this subflow was too low compared to the other subflows of this Multipath TCP connection.
- o Middlebox interference (code 0x06). Middlebox interference has been detected over this subflow making MPTCP signaling invalid. For example, this may be sent if the checksum does not validate.

3.7. Fallback

Sometimes, middleboxes will exist on a path that could prevent the operation of MPTCP. MPTCP has been designed in order to cope with many middlebox modifications (see Section 6), but there are still some cases where a subflow could fail to operate within the MPTCP requirements. These cases are notably the following: the loss of

MPTCP options on a path, and the modification of payload data. If such an event occurs, it is necessary to "fall back" to the previous, safe operation. This may be either falling back to regular TCP or removing a problematic subflow.

At the start of an MPTCP connection (i.e., the first subflow), it is important to ensure that the path is fully MPTCP capable and the necessary MPTCP options can reach each host. The handshake as described in Section 3.1 SHOULD fall back to regular TCP if either of the SYN messages do not have the MPTCP options: this is the same, and desired, behavior in the case where a host is not MPTCP capable, or the path does not support the MPTCP options. When attempting to join an existing MPTCP connection (Section 3.2), if a path is not MPTCP capable and the MPTCP options do not get through on the SYNs, the subflow will be closed according to the MP_JOIN logic.

There is, however, another corner case that should be addressed. That is one of MPTCP options getting through on the SYN, but not on regular packets. This can be resolved if the subflow is the first subflow, and thus all data in flight is contiguous, using the following rules.

A sender MUST include a DSS option with data sequence mapping in every segment until one of the sent segments has been acknowledged with a DSS option containing a Data ACK. Upon reception of the acknowledgment, the sender has the confirmation that the DSS option passes in both directions and may choose to send fewer DSS options than once per segment.

If, however, an ACK is received for data (not just for the SYN) without a DSS option containing a Data ACK, the sender determines the path is not MPTCP capable. In the case of this occurring on an additional subflow (i.e., one started with MP_JOIN), the host MUST close the subflow with a RST, which SHOULD contain a MP_TPCRST option (Section 3.6) with a "Middlebox interference" reason code.

In the case of such an ACK being received on the first subflow (i.e., that started with MP_CAPABLE), before any additional subflows are added, the implementation MUST drop out of an MPTCP mode, back to regular TCP. The sender will send one final data sequence mapping, with the Data-Level Length value of 0 indicating an infinite mapping (to inform the other end in case the path drops options in one direction only), and then revert to sending data on the single subflow without any MPTCP options.

If a subflow breaks during operation, e.g. if it is re-routed and MPTCP options are no longer permitted, then once this is detected (by the subflow-level receive buffer filling up, since there is no

mapping available in order to DATA_ACK this data), the subflow SHOULD be treated as broken and closed with a RST, since no data can be delivered to the application layer, and no fallback signal can be reliably sent. This RST SHOULD include the MP_TCPRST option (Section 3.6) with a "Middlebox interference" reason code.

These rules should cover all cases where such a failure could happen: whether it's on the forward or reverse path and whether the server or the client first sends data.

So far this section has discussed the loss of MPTCP options, either initially, or during the course of the connection. As described in Section 3.3, each portion of data for which there is a mapping is protected by a checksum, if checksums have been negotiated. This mechanism is used to detect if middleboxes have made any adjustments to the payload (added, removed, or changed data). A checksum will fail if the data has been changed in any way. This will also detect if the length of data on the subflow is increased or decreased, and this means the data sequence mapping is no longer valid. The sender no longer knows what subflow-level sequence number the receiver is genuinely operating at (the middlebox will be faking ACKs in return), and it cannot signal any further mappings. Furthermore, in addition to the possibility of payload modifications that are valid at the application layer, there is the possibility that such modifications could be triggered across MPTCP segment boundaries, corrupting the data. Therefore, all data from the start of the segment that failed the checksum onwards is not trustworthy.

Note that if checksum usage has not been negotiated, this fallback mechanism cannot be used unless there is some higher or lower layer signal to inform the MPTCP implementation that the payload has been tampered with.

When multiple subflows are in use, the data in flight on a subflow will likely involve data that is not contiguously part of the connection-level stream, since segments will be spread across the multiple subflows. Due to the problems identified above, it is not possible to determine what adjustment has done to the data (notably, any changes to the subflow sequence numbering). Therefore, it is not possible to recover the subflow, and the affected subflow must be immediately closed with a RST, featuring an MP_FAIL option (Figure 16), which defines the data sequence number at the start of the segment (defined by the data sequence mapping) that had the checksum failure. Note that the MP_FAIL option requires the use of the full 64-bit sequence number, even if 32-bit sequence numbers are normally in use in the DSS signals on the path.

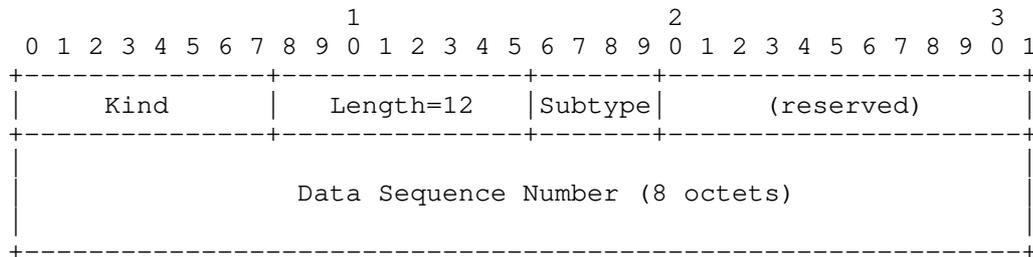


Figure 16: Fallback (MP_FAIL) Option

The receiver of this option MUST discard all data following the data sequence number specified. Failed data MUST NOT be DATA_ACKed and so will be retransmitted on other subflows (Section 3.3.6).

A special case is when there is a single subflow and it fails with a checksum error. If it is known that all unacknowledged data in flight is contiguous (which will usually be the case with a single subflow), an infinite mapping can be applied to the subflow without the need to close it first, and essentially turn off all further MPTCP signaling. In this case, if a receiver identifies a checksum failure when there is only one path, it will send back an MP_FAIL option on the subflow-level ACK, referring to the data-level sequence number of the start of the segment on which the checksum error was detected. The sender will receive this, and if all unacknowledged data in flight is contiguous, will signal an infinite mapping. This infinite mapping will be a DSS option (Section 3.3) on the first new packet, containing a data sequence mapping that acts retroactively, referring to the start of the subflow sequence number of the most recent segment that was known to be delivered intact (i.e. was successfully DATA_ACKed). From that point onwards, data can be altered by a middlebox without affecting MPTCP, as the data stream is equivalent to a regular, legacy TCP session. Whilst in theory paths may only be damaged in one direction, and the MP_FAIL signal affects only one direction of traffic, for implementation simplicity, the receiver of an MP_FAIL MUST also respond with an MP_FAIL in the reverse direction and entirely revert to a regular TCP session.

In the rare case that the data is not contiguous (which could happen when there is only one subflow but it is retransmitting data from a subflow that has recently been uncleanly closed), the receiver MUST close the subflow with a RST with MP_FAIL. The receiver MUST discard all data that follows the data sequence number specified. The sender MAY attempt to create a new subflow belonging to the same connection, and, if it chooses to do so, SHOULD place the single subflow immediately in single-path mode by setting an infinite data sequence

mapping. This mapping will begin from the data-level sequence number that was declared in the MP_FAIL.

After a sender signals an infinite mapping, it MUST only use subflow ACKs to clear its send buffer. This is because Data ACKs may become misaligned with the subflow ACKs when middleboxes insert or delete data. The receive SHOULD stop generating Data ACKs after it receives an infinite mapping.

When a connection has fallen back with an infinite mapping, only one subflow can send data; otherwise, the receiver would not know how to reorder the data. In practice, this means that all MPTCP subflows will have to be terminated except one. Once MPTCP falls back to regular TCP, it MUST NOT revert to MPTCP later in the connection.

It should be emphasized that MPTCP is not attempting to prevent the use of middleboxes that want to adjust the payload. An MPTCP-aware middlebox could provide such functionality by also rewriting checksums.

3.8. Error Handling

In addition to the fallback mechanism as described above, the standard classes of TCP errors may need to be handled in an MPTCP-specific way. Note that changing semantics -- such as the relevance of a RST -- are covered in Section 4. Where possible, we do not want to deviate from regular TCP behavior.

The following list covers possible errors and the appropriate MPTCP behavior:

- o Unknown token in MP_JOIN (or HMAC failure in MP_JOIN ACK, or missing MP_JOIN in SYN/ACK response): send RST (analogous to TCP's behavior on an unknown port)
- o DSN out of window (during normal operation): drop the data, do not send Data ACKs
- o Remove request for unknown address ID: silently ignore

3.9. Heuristics

There are a number of heuristics that are needed for performance or deployment but that are not required for protocol correctness. In this section, we detail such heuristics. Note that discussion of buffering and certain sender and receiver window behaviors are presented in Sections 3.3.4 and 3.3.5, as well as retransmission in Section 3.3.6.

3.9.1. Port Usage

Under typical operation, an MPTCP implementation SHOULD use the same ports as already in use. In other words, the destination port of a SYN containing an MP_JOIN option SHOULD be the same as the remote port of the first subflow in the connection. The local port for such SYNs SHOULD also be the same as for the first subflow (and as such, an implementation SHOULD reserve ephemeral ports across all local IP addresses), although there may be cases where this is infeasible. This strategy is intended to maximize the probability of the SYN being permitted by a firewall or NAT at the recipient and to avoid confusing any network monitoring software.

There may also be cases, however, where a host wishes to signal that a specific port should be used, and this facility is provided in the ADD_ADDR option as documented in Section 3.4.1. It is therefore feasible to allow multiple subflows between the same two addresses but using different port pairs, and such a facility could be used to allow load balancing within the network based on 5-tuples (e.g., some ECMP implementations [RFC2992]).

3.9.2. Delayed Subflow Start and Subflow Symmetry

Many TCP connections are short-lived and consist only of a few segments, and so the overheads of using MPTCP outweigh any benefits. A heuristic is required, therefore, to decide when to start using additional subflows in an MPTCP connection. Experimental deployments have shown that MPTCP can be applied in a range of scenarios so an implementation is likely to need to take into account factors including the type of traffic being sent and duration of session, and this information MAY be signalled by the application layer.

However, for standard TCP traffic, a suggested general-purpose heuristic that an implementation MAY choose to employ is as follows.

If a host has data buffered for its peer (which implies that the application has received a request for data), the host opens one subflow for each initial window's worth of data that is buffered.

Consideration should also be given to limiting the rate of adding new subflows, as well as limiting the total number of subflows open for a particular connection. A host may choose to vary these values based on its load or knowledge of traffic and path characteristics.

Note that this heuristic alone is probably insufficient. Traffic for many common applications, such as downloads, is highly asymmetric and the host that is multihomed may well be the client that will never fill its buffers, and thus never use MPTCP according to this

heuristic. Advanced APIs that allow an application to signal its traffic requirements would aid in these decisions.

An additional time-based heuristic could be applied, opening additional subflows after a given period of time has passed. This would alleviate the above issue, and also provide resilience for low-bandwidth but long-lived applications.

Another issue is that both communicating hosts may simultaneously try to set up a subflow between the same pair of addresses. This leads to an inefficient use of resources.

If the same ports are used on all subflows, as recommended above, then standard TCP simultaneous open logic should take care of this situation and only one subflow will be established between the address pairs. However, this relies on the same ports being used at both end hosts. If a host does not support TCP simultaneous open, it is RECOMMENDED that some element of randomization is applied to the time to wait before opening new subflows, so that only one subflow is created between a given address pair. If, however, hosts signal additional ports to use (for example, for leveraging ECMP on-path), this heuristic is not appropriate.

This section has shown some of the considerations that an implementer should give when developing MPTCP heuristics, but is not intended to be prescriptive.

3.9.3. Failure Handling

Requirements for MPTCP's handling of unexpected signals have been given in Section 3.8. There are other failure cases, however, where a hosts can choose appropriate behavior.

For example, Section 3.1 suggests that a host SHOULD fall back to trying regular TCP SYNs after one or more failures of MPTCP SYNs for a connection. A host may keep a system-wide cache of such information, so that it can back off from using MPTCP, firstly for that particular destination host, and eventually on a whole interface, if MPTCP connections continue failing. The duration of such a cache would be implementation-specific.

Another failure could occur when the MP_JOIN handshake fails. Section 3.8 specifies that an incorrect handshake MUST lead to the subflow being closed with a RST. A host operating an active intrusion detection system may choose to start blocking MP_JOIN packets from the source host if multiple failed MP_JOIN attempts are seen. From the connection initiator's point of view, if an MP_JOIN fails, it SHOULD NOT attempt to connect to the same IP address and

port during the lifetime of the connection, unless the other host refreshes the information with another ADD_ADDR option. Note that the ADD_ADDR option is informational only, and does not guarantee the other host will attempt a connection.

In addition, an implementation may learn, over a number of connections, that certain interfaces or destination addresses consistently fail and may default to not trying to use MPTCP for these. Behavior could also be learned for particularly badly performing subflows or subflows that regularly fail during use, in order to temporarily choose not to use these paths.

4. Semantic Issues

In order to support multipath operation, the semantics of some TCP components have changed. To aid clarity, this section collects these semantic changes as a reference.

Sequence number: The (in-header) TCP sequence number is specific to the subflow. To allow the receiver to reorder application data, an additional data-level sequence space is used. In this data-level sequence space, the initial SYN and the final DATA_FIN occupy 1 octet of sequence space. This is to ensure these signals are acknowledged at the connection level. There is an explicit mapping of data sequence space to subflow sequence space, which is signaled through TCP options in data packets.

ACK: The ACK field in the TCP header acknowledges only the subflow sequence number, not the data-level sequence space. Implementations SHOULD NOT attempt to infer a data-level acknowledgment from the subflow ACKs. This separates subflow- and connection-level processing at an end host.

Duplicate ACK: A duplicate ACK that includes any MPTCP signaling (with the exception of the DSS option) MUST NOT be treated as a signal of congestion. To limit the chances of non-MPTCP-aware entities mistakenly interpreting duplicate ACKs as a signal of congestion, MPTCP SHOULD NOT send more than two duplicate ACKs containing (non-DSS) MPTCP signals in a row.

Receive Window: The receive window in the TCP header indicates the amount of free buffer space for the whole data-level connection (as opposed to for this subflow) that is available at the receiver. This is the same semantics as regular TCP, but to maintain these semantics the receive window must be interpreted at the sender as relative to the sequence number given in the DATA_ACK rather than the subflow ACK in the TCP header. In this way, the original flow control role is preserved. Note that some

middleboxes may change the receive window, and so a host SHOULD use the maximum value of those recently seen on the constituent subflows for the connection-level receive window, and also needs to maintain a subflow-level window for subflow-level processing.

FIN: The FIN flag in the TCP header applies only to the subflow it is sent on, not to the whole connection. For connection-level FIN semantics, the DATA_FIN option is used.

RST: The RST flag in the TCP header applies only to the subflow it is sent on, not to the whole connection. The MP_FASTCLOSE option provides the fast close functionality of a RST at the MPTCP connection level.

Address List: Address list management (i.e., knowledge of the local and remote hosts' lists of available IP addresses) is handled on a per-connection basis (as opposed to per subflow, per host, or per pair of communicating hosts). This permits the application of per-connection local policy. Adding an address to one connection (either explicitly through an Add Address message, or implicitly through a Join) has no implication for other connections between the same pair of hosts.

5-tuple: The 5-tuple (protocol, local address, local port, remote address, remote port) presented by kernel APIs to the application layer in a non-multipath-aware application is that of the first subflow, even if the subflow has since been closed and removed from the connection. This decision, and other related API issues, are discussed in more detail in [RFC6897].

5. Security Considerations

As identified in [RFC6181], the addition of multipath capability to TCP will bring with it a number of new classes of threat. In order to prevent these, [RFC6182] presents a set of requirements for a security solution for MPTCP. The fundamental goal is for the security of MPTCP to be "no worse" than regular TCP today, and the key security requirements are:

- o Provide a mechanism to confirm that the parties in a subflow handshake are the same as in the original connection setup.
- o Provide verification that the peer can receive traffic at a new address before using it as part of a connection.
- o Provide replay protection, i.e., ensure that a request to add/remove a subflow is 'fresh'.

In order to achieve these goals, MPTCP includes a hash-based handshake algorithm documented in Sections 3.1 and 3.2.

The security of the MPTCP connection hangs on the use of keys that are shared once at the start of the first subflow, and are never sent again over the network (unless used in the fast close mechanism, Section 3.5). To ease demultiplexing while not giving away any cryptographic material, future subflows use a truncated cryptographic hash of this key as the connection identification "token". The keys are concatenated and used as keys for creating Hash-based Message Authentication Codes (HMACs) used on subflow setup, in order to verify that the parties in the handshake are the same as in the original connection setup. It also provides verification that the peer can receive traffic at this new address. Replay attacks would still be possible when only keys are used; therefore, the handshakes use single-use random numbers (nonces) at both ends -- this ensures the HMAC will never be the same on two handshakes. Guidance on generating random numbers suitable for use as keys is given in [RFC4086] and discussed in Section 3.1. The nonces are valid for the lifetime of the TCP connection attempt. HMAC is also used to secure the ADD_ADDR option, due to the threats identified in [RFC7430].

The use of crypto capability bits in the initial connection handshake to negotiate use of a particular algorithm allows the deployment of additional crypto mechanisms in the future. This negotiation would nevertheless be susceptible to a bid-down attack by an on-path active attacker who could modify the crypto capability bits in the response from the receiver to use a less secure crypto mechanism. The security mechanism presented in this document should therefore protect against all forms of flooding and hijacking attacks discussed in [RFC6181].

The version negotiation specified in Section 3.1, if differing MPTCP versions shared a common negotiation format, would allow an on-path attacker to apply a theoretical bid-down attack. Since the v1 and v0 protocols have a different handshake, such an attack would require the client to re-establish the connection using v0, and this being supported by the server. Note that an on-path attacker would have access to the raw data, negating any other TCP-level security mechanisms. Also a change from RFC6824 has removed the subflow identifier from the MP_PRIO option (Section 3.3.8), to remove the theoretical attack where a subflow could be placed in "backup" mode by an attacker.

During normal operation, regular TCP protection mechanisms (such as ensuring sequence numbers are in-window) will provide the same level of protection against attacks on individual TCP subflows as exists for regular TCP today. Implementations will introduce additional

buffers compared to regular TCP, to reassemble data at the connection level. The application of window sizing will minimize the risk of denial-of-service attacks consuming resources.

As discussed in Section 3.4.1, a host may advertise its private addresses, but these might point to different hosts in the receiver's network. The MP_JOIN handshake (Section 3.2) will ensure that this does not succeed in setting up a subflow to the incorrect host. However, it could still create unwanted TCP handshake traffic. This feature of MPTCP could be a target for denial-of-service exploits, with malicious participants in MPTCP connections encouraging the recipient to target other hosts in the network. Therefore, implementations should consider heuristics (Section 3.9) at both the sender and receiver to reduce the impact of this.

To further protect against malicious ADD_ADDR messages sent by an off-path attacker, the ADD_ADDR includes an HMAC using the keys negotiated during the handshake. This effectively prevents an attacker from diverting an MPTCP connection through an off-path ADD_ADDR injection into the stream.

A small security risk could theoretically exist with key reuse, but in order to accomplish a replay attack, both the sender and receiver keys, and the sender and receiver random numbers, in the MP_JOIN handshake (Section 3.2) would have to match.

Whilst this specification defines a "medium" security solution, meeting the criteria specified at the start of this section and the threat analysis ([RFC6181]), since attacks only ever get worse, it is likely that a future version of MPTCP would need to be able to support stronger security. There are several ways the security of MPTCP could potentially be improved; some of these would be compatible with MPTCP as defined in this document, whilst others may not be. For now, the best approach is to get experience with the current approach, establish what might work, and check that the threat analysis is still accurate.

Possible ways of improving MPTCP security could include:

- o defining a new MPTCP cryptographic algorithm, as negotiated in MP_CAPABLE. A sub-case could be to include an additional deployment assumption, such as stateful servers, in order to allow a more powerful algorithm to be used.
- o defining how to secure data transfer with MPTCP, whilst not changing the signaling part of the protocol.

- o defining security that requires more option space, perhaps in conjunction with a "long options" proposal for extending the TCP options space (such as those surveyed in [TCPL0]), or perhaps building on the current approach with a second stage of MPTCP-option-based security.
- o revisiting the working group's decision to exclusively use TCP options for MPTCP signaling, and instead look at also making use of the TCP payloads.

MPTCP has been designed with several methods available to indicate a new security mechanism, including:

- o available flags in MP_CAPABLE (Figure 4);
- o available subtypes in the MPTCP option (Figure 3);
- o the version field in MP_CAPABLE (Figure 4);

6. Interactions with Middleboxes

Multipath TCP was designed to be deployable in the present world. Its design takes into account "reasonable" existing middlebox behavior. In this section, we outline a few representative middlebox-related failure scenarios and show how Multipath TCP handles them. Next, we list the design decisions multipath has made to accommodate the different middleboxes.

A primary concern is our use of a new TCP option. Middleboxes should forward packets with unknown options unchanged, yet there are some that don't. These we expect will either strip options and pass the data, drop packets with new options, copy the same option into multiple segments (e.g., when doing segmentation), or drop options during segment coalescing.

MPTCP uses a single new TCP option "Kind", and all message types are defined by "subtype" values (see Section 8). This should reduce the chances of only some types of MPTCP options being passed, and instead the key differing characteristics are different paths, and the presence of the SYN flag.

MPTCP SYN packets on the first subflow of a connection contain the MP_CAPABLE option (Section 3.1). If this is dropped, MPTCP SHOULD fall back to regular TCP. If packets with the MP_JOIN option (Section 3.2) are dropped, the paths will simply not be used.

If a middlebox strips options but otherwise passes the packets unchanged, MPTCP will behave safely. If an MP_CAPABLE option is

dropped on either the outgoing or the return path, the initiating host can fall back to regular TCP, as illustrated in Figure 17 and discussed in Section 3.1.

Subflow SYNs contain the MP_JOIN option. If this option is stripped on the outgoing path, the SYN will appear to be a regular SYN to Host B. Depending on whether there is a listening socket on the target port, Host B will reply either with SYN/ACK or RST (subflow connection fails). When Host A receives the SYN/ACK it sends a RST because the SYN/ACK does not contain the MP_JOIN option and its token. Either way, the subflow setup fails, but otherwise does not affect the MPTCP connection as a whole.

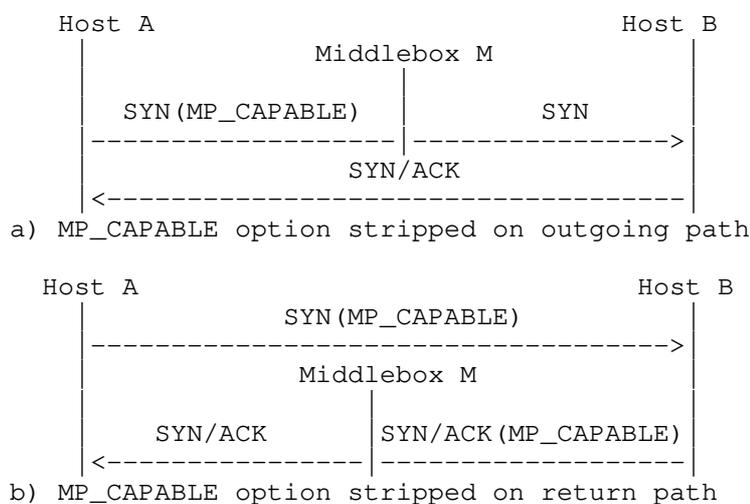


Figure 17: Connection Setup with Middleboxes that Strip Options from Packets

We now examine data flow with MPTCP, assuming the flow is correctly set up, which implies the options in the SYN packets were allowed through by the relevant middleboxes. If options are allowed through and there is no resegmentation or coalescing to TCP segments, Multipath TCP flows can proceed without problems.

The case when options get stripped on data packets has been discussed in the Fallback section. If only some MPTCP options are stripped, behavior is not deterministic. If some data sequence mappings are lost, the connection can continue so long as mappings exist for the subflow-level data (e.g., if multiple maps have been sent that reinforce each other). If some subflow-level space is left unmapped, however, the subflow is treated as broken and is closed, through the process described in Section 3.7. MPTCP should survive with a loss

of some Data ACKs, but performance will degrade as the fraction of stripped options increases. We do not expect such cases to appear in practice, though: most middleboxes will either strip all options or let them all through.

We end this section with a list of middlebox classes, their behavior, and the elements in the MPTCP design that allow operation through such middleboxes. Issues surrounding dropping packets with options or stripping options were discussed above, and are not included here:

- o NATs [RFC3022] (Network Address (and Port) Translators) change the source address (and often source port) of packets. This means that a host will not know its public-facing address for signaling in MPTCP. Therefore, MPTCP permits implicit address addition via the MP_JOIN option, and the handshake mechanism ensures that connection attempts to private addresses [RFC1918], since they are authenticated, will only set up subflows to the correct hosts. Explicit address removal is undertaken by an Address ID to allow no knowledge of the source address.
- o Performance Enhancing Proxies (PEPs) [RFC3135] might proactively ACK data to increase performance. MPTCP, however, relies on accurate congestion control signals from the end host, and non-MPTCP-aware PEPs will not be able to provide such signals. MPTCP will, therefore, fall back to single-path TCP, or close the problematic subflow (see Section 3.7).
- o Traffic Normalizers [norm] may not allow holes in sequence numbers, and may cache packets and retransmit the same data. MPTCP looks like standard TCP on the wire, and will not retransmit different data on the same subflow sequence number. In the event of a retransmission, the same data will be retransmitted on the original TCP subflow even if it is additionally retransmitted at the connection level on a different subflow.
- o Firewalls [RFC2979] might perform initial sequence number randomization on TCP connections. MPTCP uses relative sequence numbers in data sequence mapping to cope with this. Like NATs, firewalls will not permit many incoming connections, so MPTCP supports address signaling (ADD_ADDR) so that a multiaddressed host can invite its peer behind the firewall/NAT to connect out to its additional interface.
- o Intrusion Detection/Prevention Systems (IDS/IPS) observe packet streams for patterns and content that could threaten a network. MPTCP may require the instrumentation of additional paths, and an MPTCP-aware IDS/IPS would need to read MPTCP tokens to correlate data from multiple subflows to maintain comparable visibility into

all of the traffic between devices. Without such changes, an IDS would get an incomplete view of the traffic, increasing the risk of missing traffic of interest (false negatives), and increasing the chances of erroneously identifying a subflow as a risk due to only seeing partial data (false positives).

- o Application-level middleboxes such as content-aware firewalls may alter the payload within a subflow, such as rewriting URIs in HTTP traffic. MPTCP will detect these using the checksum and close the affected subflow(s), if there are other subflows that can be used. If all subflows are affected, multipath will fall back to TCP, allowing such middleboxes to change the payload. MPTCP-aware middleboxes should be able to adjust the payload and MPTCP metadata in order not to break the connection.

In addition, all classes of middleboxes may affect TCP traffic in the following ways:

- o TCP options may be removed, or packets with unknown options dropped, by many classes of middleboxes. It is intended that the initial SYN exchange, with a TCP option, will be sufficient to identify the path capabilities. If such a packet does not get through, MPTCP will end up falling back to regular TCP.
- o Segmentation/Coalescing (e.g., TCP segmentation offloading) might copy options between packets and might strip some options. MPTCP's data sequence mapping includes the relative subflow sequence number instead of using the sequence number in the segment. In this way, the mapping is independent of the packets that carry it.
- o The receive window may be shrunk by some middleboxes at the subflow level. MPTCP will use the maximum window at data level, but will also obey subflow-specific windows.

7. Acknowledgments

The authors gratefully acknowledge significant input into this document from Sebastien Barre and Andrew McDonald.

The authors also wish to acknowledge reviews and contributions from Iljitsch van Beijnum, Lars Eggert, Marcelo Bagnulo, Robert Hancock, Pasi Sarolahti, Toby Moncaster, Philip Eardley, Sergio Lembo, Lawrence Conroy, Yoshifumi Nishida, Bob Briscoe, Stein Gjessing, Andrew McGregor, Georg Hampel, Anumita Biswas, Wes Eddy, Alexey Melnikov, Francis Dupont, Adrian Farrel, Barry Leiba, Robert Sparks, Sean Turner, Stephen Farrell, Martin Stiernerling, Gregory Detal, Fabien Duchene, Xavier de Foy, Rahul Jadhav, Klemens Schragel, Mirja

Kuehlewind, Sheng Jiang, Alissa Cooper, Ines Robles, Roman Danyliw, Adam Roach, Barry Leiba, Alexey Melnikov, Eric Vyncke, and Ben Kaduk.

8. IANA Considerations

This document obsoletes RFC6824 and as such IANA is requested to update the TCP option space registry to point to this document for Multipath TCP, as follows:

Kind	Length	Meaning	Reference
30	N	Multipath TCP (MPTCP)	This document

Table 1: TCP Option Kind Numbers

8.1. MPTCP Option Subtypes

The 4-bit MPTCP subtype sub-registry ("MPTCP Option Subtypes" under the "Transmission Control Protocol (TCP) Parameters" registry) was defined in RFC6824. Since RFC6824 was an Experimental not Standards Track RFC, and since no further entries have occurred beyond those pointing to RFC6824, IANA is requested to replace the existing registry with Table 2 and with the following explanatory note.

Note: This registry specifies the MPTCP Option Subtypes for MPTCP v1, which obsoletes the Experimental MPTCP v0. For the MPTCP v0 subtypes, please refer to RFC6824.

Value	Symbol	Name	Reference
0x0	MP_CAPABLE	Multipath Capable	This document, Section 3.1
0x1	MP_JOIN	Join Connection	This document, Section 3.2
0x2	DSS	Data Sequence Signal (Data ACK and data sequence mapping)	This document, Section 3.3
0x3	ADD_ADDR	Add Address	This document, Section 3.4.1
0x4	REMOVE_ADDR	Remove Address	This document, Section 3.4.2
0x5	MP_PRIO	Change Subflow Priority	This document, Section 3.3.8
0x6	MP_FAIL	Fallback	This document, Section 3.7
0x7	MP_FASTCLOSE	Fast Close	This document, Section 3.5
0x8	MP_TCRST	Subflow Reset	This document, Section 3.6
0xf	MP_EXPERIMENTAL	Reserved for private experiments	

Table 2: MPTCP Option Subtypes

Values 0x9 through 0xe are currently unassigned. Option 0xf is reserved for use by private experiments. Its use may be formalized in a future specification. Future assignments in this registry are to be defined by Standards Action as defined by [RFC8126]. Assignments consist of the MPTCP subtype's symbolic name and its associated value, and a reference to its specification.

8.2. MPTCP Handshake Algorithms

The "MPTCP Handshake Algorithms" sub-registry under the "Transmission Control Protocol (TCP) Parameters" registry was defined in RFC6824. Since RFC6824 was an Experimental not Standards Track RFC, and since

no further entries have occurred beyond those pointing to RFC6824, IANA is requested to replace the existing registry with Table 3 and with the following explanatory note.

Note: This registry specifies the MPTCP Handshake Algorithms for MPTCP v1, which obsoletes the Experimental MPTCP v0. For the MPTCP v0 subtypes, please refer to RFC6824.

Flag Bit	Meaning	Reference
A	Checksum required	This document, Section 3.1
B	Extensibility	This document, Section 3.1
C	Do not attempt to establish new subflows to the source address.	This document, Section 3.1
D-G	Unassigned	
H	HMAC-SHA256	This document, Section 3.2

Table 3: MPTCP Handshake Algorithms

Note that the meanings of bits D through H can be dependent upon bit B, depending on how Extensibility is defined in future specifications; see Section 3.1 for more information.

Future assignments in this registry are also to be defined by Standards Action as defined by [RFC8126]. Assignments consist of the value of the flags, a symbolic name for the algorithm, and a reference to its specification.

8.3. MP_TCPRST Reason Codes

IANA is requested to create a further sub-registry, "MPTCP MP_TCPRST Reason Codes" under the "Transmission Control Protocol (TCP) Parameters" registry, based on the reason code in MP_TCPRST (Section 3.6) message. Initial values for this registry are given in Table 4; future assignments are to be defined by Specification Required as defined by [RFC8126]. Assignments consist of the value of the code, a short description of its meaning, and a reference to its specification. The maximum value is 0xff.

As guidance to the Designated Expert [RFC8126], assignments should not normally be refused unless codepoint space is becoming scarce, providing that there is a clear distinction from other, already-

existing codes, and also providing there is sufficient guidance for implementors both sending and receiving these codes.

Code	Meaning	Reference
0x00	Unspecified TCP error	This document, Section 3.6
0x01	MPTCP specific error	This document, Section 3.6
0x02	Lack of resources	This document, Section 3.6
0x03	Administratively prohibited	This document, Section 3.6
0x04	Too much outstanding data	This document, Section 3.6
0x05	Unacceptable performance	This document, Section 3.6
0x06	Middlebox interference	This document, Section 3.6

Table 4: MPTCP MP_TCP_RST Reason Codes

9. References

9.1. Normative References

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <<https://www.rfc-editor.org/info/rfc793>>.
- [RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, DOI 10.17487/RFC2104, February 1997, <<https://www.rfc-editor.org/info/rfc2104>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5961] Ramaiah, A., Stewart, R., and M. Dalal, "Improving TCP's Robustness to Blind In-Window Attacks", RFC 5961, DOI 10.17487/RFC5961, August 2010, <<https://www.rfc-editor.org/info/rfc5961>>.
- [RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <<https://www.rfc-editor.org/info/rfc6234>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

9.2. Informative References

- [deployments] Bonaventure, O. and S. Seo, "Multipath TCP Deployments", IETF Journal 2016, November 2016, <<https://www.ietfjournal.org/multipath-tcp-deployments/>>.
- [howhard] Raiciu, C., Paasch, C., Barre, S., Ford, A., Honda, M., Duchene, F., Bonaventure, O., and M. Handley, "How Hard Can It Be? Designing and Implementing a Deployable Multipath TCP", Usenix Symposium on Networked Systems Design and Implementation 2012, 2012, <<https://www.usenix.org/conference/nsdi12/how-hard-can-it-be-designing-and-implementing-deployable-multipath-tcp>>.
- [norm] Handley, M., Paxson, V., and C. Kreibich, "Network Intrusion Detection: Evasion, Traffic Normalization, and End-to-End Protocol Semantics", Usenix Security 2001, 2001, <http://www.usenix.org/events/sec01/full_papers/handley/handley.pdf>.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <<https://www.rfc-editor.org/info/rfc1122>>.
- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <<https://www.rfc-editor.org/info/rfc1918>>.
- [RFC2018] Mathis, M., Mahdavi, J., Floyd, S., and A. Romanow, "TCP Selective Acknowledgment Options", RFC 2018, DOI 10.17487/RFC2018, October 1996, <<https://www.rfc-editor.org/info/rfc2018>>.
- [RFC2979] Freed, N., "Behavior of and Requirements for Internet Firewalls", RFC 2979, DOI 10.17487/RFC2979, October 2000, <<https://www.rfc-editor.org/info/rfc2979>>.
- [RFC2992] Hopps, C., "Analysis of an Equal-Cost Multi-Path Algorithm", RFC 2992, DOI 10.17487/RFC2992, November 2000, <<https://www.rfc-editor.org/info/rfc2992>>.

- [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, DOI 10.17487/RFC3022, January 2001, <<https://www.rfc-editor.org/info/rfc3022>>.
- [RFC3135] Border, J., Kojo, M., Griner, J., Montenegro, G., and Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations", RFC 3135, DOI 10.17487/RFC3135, June 2001, <<https://www.rfc-editor.org/info/rfc3135>>.
- [RFC4086] Eastlake 3rd, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", BCP 106, RFC 4086, DOI 10.17487/RFC4086, June 2005, <<https://www.rfc-editor.org/info/rfc4086>>.
- [RFC4987] Eddy, W., "TCP SYN Flooding Attacks and Common Mitigations", RFC 4987, DOI 10.17487/RFC4987, August 2007, <<https://www.rfc-editor.org/info/rfc4987>>.
- [RFC5681] Allman, M., Paxson, V., and E. Blanton, "TCP Congestion Control", RFC 5681, DOI 10.17487/RFC5681, September 2009, <<https://www.rfc-editor.org/info/rfc5681>>.
- [RFC6181] Bagnulo, M., "Threat Analysis for TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6181, DOI 10.17487/RFC6181, March 2011, <<https://www.rfc-editor.org/info/rfc6181>>.
- [RFC6182] Ford, A., Raiciu, C., Handley, M., Barre, S., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development", RFC 6182, DOI 10.17487/RFC6182, March 2011, <<https://www.rfc-editor.org/info/rfc6182>>.
- [RFC6356] Raiciu, C., Handley, M., and D. Wischik, "Coupled Congestion Control for Multipath Transport Protocols", RFC 6356, DOI 10.17487/RFC6356, October 2011, <<https://www.rfc-editor.org/info/rfc6356>>.
- [RFC6528] Gont, F. and S. Bellovin, "Defending against Sequence Number Attacks", RFC 6528, DOI 10.17487/RFC6528, February 2012, <<https://www.rfc-editor.org/info/rfc6528>>.
- [RFC6897] Scharf, M. and A. Ford, "Multipath TCP (MPTCP) Application Interface Considerations", RFC 6897, DOI 10.17487/RFC6897, March 2013, <<https://www.rfc-editor.org/info/rfc6897>>.

- [RFC7323] Borman, D., Braden, B., Jacobson, V., and R. Scheffenegger, Ed., "TCP Extensions for High Performance", RFC 7323, DOI 10.17487/RFC7323, September 2014, <<https://www.rfc-editor.org/info/rfc7323>>.
- [RFC7413] Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, DOI 10.17487/RFC7413, December 2014, <<https://www.rfc-editor.org/info/rfc7413>>.
- [RFC7430] Bagnulo, M., Paasch, C., Gont, F., Bonaventure, O., and C. Raiciu, "Analysis of Residual Threats and Possible Fixes for Multipath TCP (MPTCP)", RFC 7430, DOI 10.17487/RFC7430, July 2015, <<https://www.rfc-editor.org/info/rfc7430>>.
- [RFC8041] Bonaventure, O., Paasch, C., and G. Detal, "Use Cases and Operational Experience with Multipath TCP", RFC 8041, DOI 10.17487/RFC8041, January 2017, <<https://www.rfc-editor.org/info/rfc8041>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [TCPLO] Ramaiah, A., "TCP option space extension", Work in Progress, March 2012.

Appendix A. Notes on Use of TCP Options

The TCP option space is limited due to the length of the Data Offset field in the TCP header (4 bits), which defines the TCP header length in 32-bit words. With the standard TCP header being 20 bytes, this leaves a maximum of 40 bytes for options, and many of these may already be used by options such as timestamp and SACK.

We have performed a brief study on the commonly used TCP options in SYN, data, and pure ACK packets, and found that there is enough room to fit all the options we propose using in this document.

SYN packets typically include Maximum Segment Size (MSS) (4 bytes), window scale (3 bytes), SACK permitted (2 bytes), and timestamp (10 bytes) options. Together these sum to 19 bytes. Some operating systems appear to pad each option up to a word boundary, thus using 24 bytes (a brief survey suggests Windows XP and Mac OS X do this, whereas Linux does not). Optimistically, therefore, we have 21 bytes spare, or 16 if it has to be word-aligned. In either case, however, the SYN versions of Multipath Capable (12 bytes) and Join (12 or 16 bytes) options will fit in this remaining space.

Note that due to the use of a 64-bit data-level sequence space, it is feasible that MPTCP will not require the timestamp option for protection against wrapped sequence numbers (PAWS [RFC7323]), since the data-level sequence space has far less chance of wrapping. Confirmation of the validity of this optimisation is for further study.

TCP data packets typically carry timestamp options in every packet, taking 10 bytes (or 12 with padding). That leaves 30 bytes (or 28, if word-aligned). The Data Sequence Signal (DSS) option varies in length depending on whether the data sequence mapping and DATA_ACK are included, and whether the sequence numbers in use are 4 or 8 octets. The maximum size of the DSS option is 28 bytes, so even that will fit in the available space. But unless a connection is both bidirectional and high-bandwidth, it is unlikely that all that option space will be required on each DSS option.

Within the DSS option, it is not necessary to include the data sequence mapping and DATA_ACK in each packet, and in many cases it may be possible to alternate their presence (so long as the mapping covers the data being sent in the following packet). It would also be possible to alternate between 4- and 8-byte sequence numbers in each option.

On subflow and connection setup, an MPTCP option is also set on the third packet (an ACK). These are 20 bytes (for Multipath Capable)

and 24 bytes (for Join), both of which will fit in the available option space.

Pure ACKs in TCP typically contain only timestamps (10 bytes). Here, Multipath TCP typically needs to encode only the DATA_ACK (maximum of 12 bytes). Occasionally, ACKs will contain SACK information. Depending on the number of lost packets, SACK may utilize the entire option space. If a DATA_ACK had to be included, then it is probably necessary to reduce the number of SACK blocks to accommodate the DATA_ACK. However, the presence of the DATA_ACK is unlikely to be necessary in a case where SACK is in use, since until at least some of the SACK blocks have been retransmitted, the cumulative data-level ACK will not be moving forward (or if it does, due to retransmissions on another path, then that path can also be used to transmit the new DATA_ACK).

The ADD_ADDR option can be between 16 and 30 bytes, depending on whether IPv4 or IPv6 is used, and whether or not the port number is present. It is unlikely that such signaling would fit in a data packet (although if there is space, it is fine to include it). It is recommended to use duplicate ACKs with no other payload or options in order to transmit these rare signals. Note this is the reason for mandating that duplicate ACKs with MPTCP options are not taken as a signal of congestion.

Appendix B. TCP Fast Open and MPTCP

TCP Fast Open (TFO) is an experimental TCP extension, described in [RFC7413], which has been introduced to allow sending data one RTT earlier than with regular TCP. This is considered a valuable gain as very short connections are very common, especially for HTTP request/response schemes. It achieves this by sending the SYN-segment together with the application's data and allowing the listener to reply immediately with data after the SYN/ACK. [RFC7413] secures this mechanism, by using a new TCP option that includes a cookie which is negotiated in a preceding connection.

When using TCP Fast Open in conjunction with MPTCP, there are two key points to take into account, detailed hereafter.

B.1. TFO cookie request with MPTCP

When a TFO initiator first connects to a listener, it cannot immediately include data in the SYN for security reasons [RFC7413]. Instead, it requests a cookie that will be used in subsequent connections. This is done with the TCP cookie request/response options, of respectively 2 bytes and 6-18 bytes (depending on the chosen cookie length).

TFO and MPTCP can be combined provided that the total length of all the options does not exceed the maximum 40 bytes possible in TCP:

- o In the SYN: MPTCP uses a 4-bytes long MP_CAPABLE option. The MPTCP and TFO options sum up to 6 bytes. With typical TCP-options using up to 19 bytes in the SYN (24 bytes if options are padded at a word boundary), there is enough space to combine the MP_CAPABLE with the TFO Cookie Request.
- o In the SYN+ACK: MPTCP uses a 12-bytes long MP_CAPABLE option, but now TFO can be as long as 18 bytes. Since the maximum option length may be exceeded, it is up to the listener to solve this by using a shorter cookie. As an example, if we consider that 19 bytes are used for classical TCP options, the maximum possible cookie length would be of 7 bytes. Note that the same limitation applies to subsequent connections, for the SYN packet (because the initiator then echoes back the cookie to the listener). Finally, if the security impact of reducing the cookie size is not deemed acceptable, the listener can reduce the amount of other TCP-options by omitting the TCP timestamps (as outlined in Appendix A).

B.2. Data sequence mapping under TFO

MPTCP uses, in the TCP establishment phase, a key exchange that is used to generate the Initial Data Sequence Numbers (IDSNs). In particular, the SYN with MP_CAPABLE occupies the first octet of the data sequence space. With TFO, one way to handle the data sent together with the SYN would be to consider an implicit DSS mapping that covers that SYN segment (since there is not enough space in the SYN to include a DSS option). The problem with that approach is that if a middlebox modifies the TFO data, this will not be noticed by MPTCP because of the absence of a DSS-checksum. For example, a TCP (but not MPTCP)-aware middlebox could insert bytes at the beginning of the stream and adapt the TCP checksum and sequence numbers accordingly. With an implicit mapping, this would give to initiator and listener a different view on the DSS-mapping, with no way to detect this inconsistency as the DSS checksum is not present.

To solve this, the TFO data must not be considered part of the Data Sequence Number space: the SYN with MP_CAPABLE still occupies the first octet of data sequence space, but then the first non-TFO data byte occupies the second octet. This guarantees that, if the use of DSS-checksum is negotiated, all data in the data sequence number space is checksummed. We also note that this does not entail a loss of functionality, because TFO-data is always only sent on the initial subflow before any attempt to create additional subflows.

B.3. Connection establishment examples

The following shows a few examples of possible TFO+MPTCP establishment scenarios.

Before an initiator can send data together with the SYN, it must request a cookie to the listener, as shown in Figure 18. This is done by simply combining the TFO and MPTCP options.

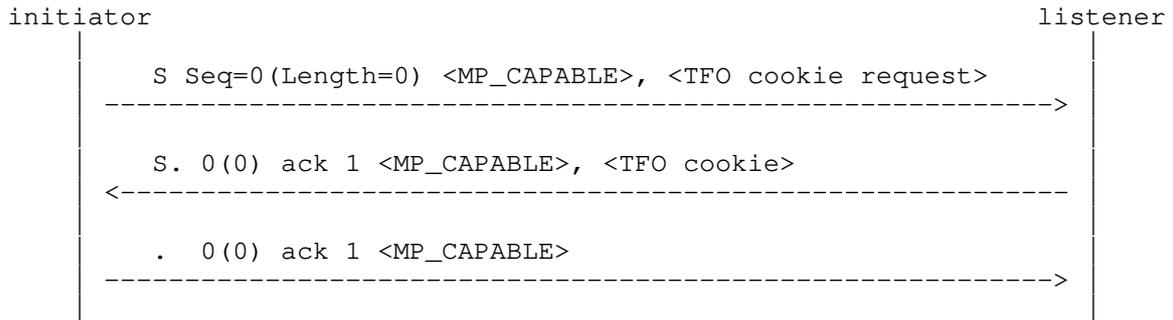


Figure 18: Cookie request - sequence number and length are annotated as Seq(Length) and used hereafter in the figures.

Once this is done, the received cookie can be used for TFO, as shown in Figure 19. In this example, the initiator first sends 20 bytes in the SYN. The listener immediately replies with 100 bytes following the SYN-ACK upon which the initiator replies with 20 more bytes. Note that the last segment in the figure has a TCP sequence number of 21, while the DSS subflow sequence number is 1 (because the TFO data is not part of the data sequence number space, as explained in Section Appendix B.2).

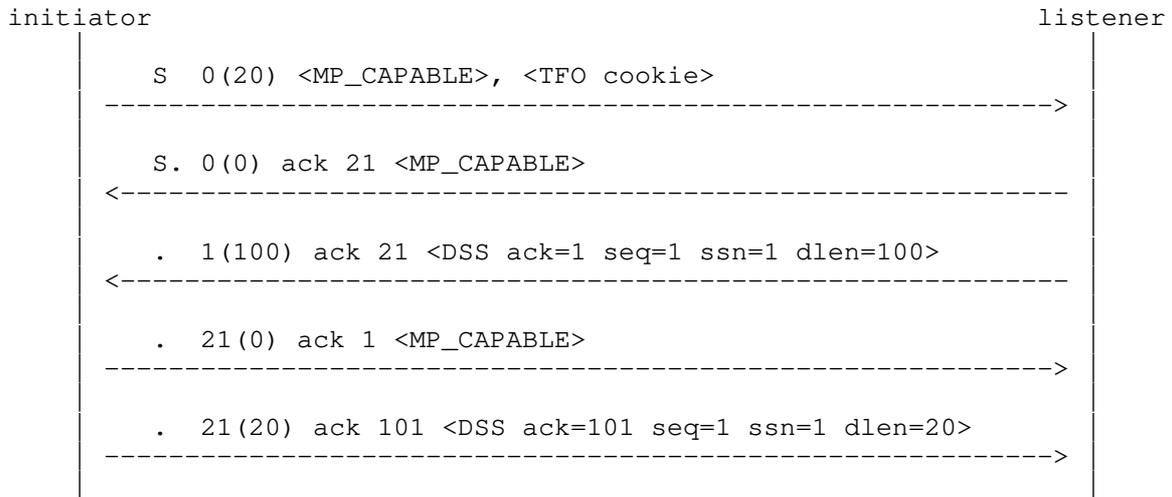


Figure 19: The listener supports TFO

In Figure 20, the listener does not support TFO. The initiator detects that no state is created in the listener (as no data is acked), and now sends the MP_CAPABLE in the third ack, in order for the listener to build its MPTCP context at then end of the establishment. Now, the tfo data, retransmitted, becomes part of the data sequence mapping because it is effectively sent (in fact re-sent) after the establishment.

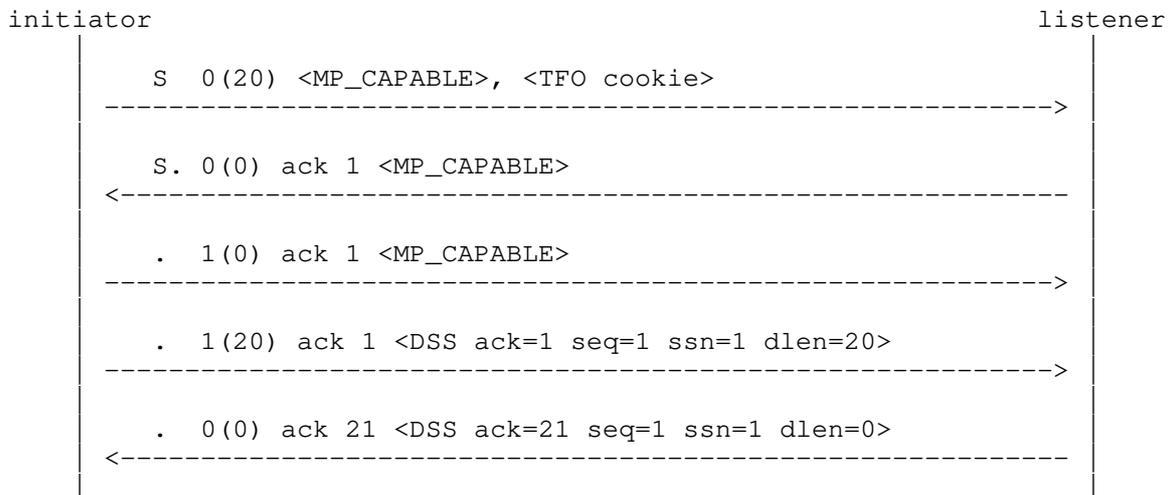


Figure 20: The listener does not support TFO

It is also possible that the listener acknowledges only part of the TFO data, as illustrated in Figure 21. The initiator will simply retransmit the missing data together with a DSS-mapping.

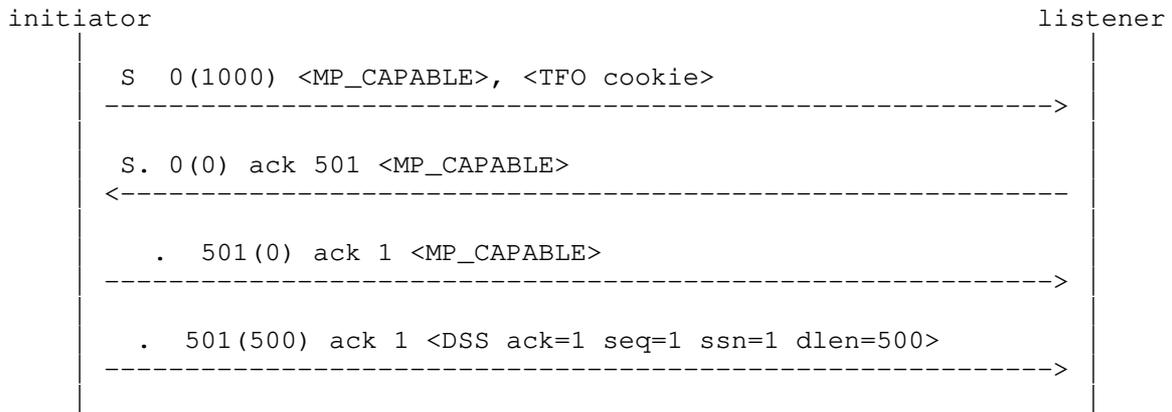


Figure 21: Partial data acknowledgement

Appendix C. Control Blocks

Conceptually, an MPTCP connection can be represented as an MPTCP protocol control block (PCB) that contains several variables that track the progress and the state of the MPTCP connection and a set of linked TCP control blocks that correspond to the subflows that have been established.

RFC 793 [RFC0793] specifies several state variables. Whenever possible, we reuse the same terminology as RFC 793 to describe the state variables that are maintained by MPTCP.

C.1. MPTCP Control Block

The MPTCP control block contains the following variable per connection.

C.1.1. Authentication and Metadata

Local.Token (32 bits): This is the token chosen by the local host on this MPTCP connection. The token must be unique among all established MPTCP connections, and is generated from the local key.

Local.Key (64 bits): This is the key sent by the local host on this MPTCP connection.

Remote.Token (32 bits): This is the token chosen by the remote host on this MPTCP connection, generated from the remote key.

Remote.Key (64 bits): This is the key chosen by the remote host on this MPTCP connection

MPTCP.Checksum (flag): This flag is set to true if at least one of the hosts has set the A bit in the MP_CAPABLE options exchanged during connection establishment, and is set to false otherwise. If this flag is set, the checksum must be computed in all DSS options.

C.1.2. Sending Side

SND.UNA (64 bits): This is the data sequence number of the next byte to be acknowledged, at the MPTCP connection level. This variable is updated upon reception of a DSS option containing a DATA_ACK.

SND.NXT (64 bits): This is the data sequence number of the next byte to be sent. SND.NXT is used to determine the value of the DSN in the DSS option.

SND.WND (32 bits with RFC 7323, 16 bits otherwise): This is the sending window. MPTCP maintains the sending window at the MPTCP connection level and the same window is shared by all subflows. All subflows use the MPTCP connection level SND.WND to compute the SEQ.WND value that is sent in each transmitted segment.

C.1.3. Receiving Side

RCV.NXT (64 bits): This is the data sequence number of the next byte that is expected on the MPTCP connection. This state variable is modified upon reception of in-order data. The value of RCV.NXT is used to specify the DATA_ACK that is sent in the DSS option on all subflows.

RCV.WND (32 bits with RFC 7323, 16 bits otherwise): This is the connection-level receive window, which is the maximum of the RCV.WND on all the subflows.

C.2. TCP Control Blocks

The MPTCP control block also contains a list of the TCP control blocks that are associated with the MPTCP connection.

Note that the TCP control block on the TCP subflows does not contain the RCV.WND and SND.WND state variables as these are maintained at the MPTCP connection level and not at the subflow level.

Inside each TCP control block, the following state variables are defined.

C.2.1. Sending Side

SND.UNA (32 bits): This is the sequence number of the next byte to be acknowledged on the subflow. This variable is updated upon reception of each TCP acknowledgment on the subflow.

SND.NXT (32 bits): This is the sequence number of the next byte to be sent on the subflow. SND.NXT is used to set the value of SEG.SEQ upon transmission of the next segment.

C.2.2. Receiving Side

RCV.NXT (32 bits): This is the sequence number of the next byte that is expected on the subflow. This state variable is modified upon reception of in-order segments. The value of RCV.NXT is copied to the SEG.ACK field of the next segments transmitted on the subflow.

RCV.WND (32 bits with RFC 7323, 16 bits otherwise): This is the subflow-level receive window that is updated with the window field from the segments received on this subflow.

Appendix D. Finite State Machine

The diagram in Figure 22 shows the Finite State Machine for connection-level closure. This illustrates how the DATA_FIN connection-level signal (indicated in the diagram as the DFIN flag on a DATA_ACK) interacts with subflow-level FINs, and permits "break-before-make" handover between subflows.

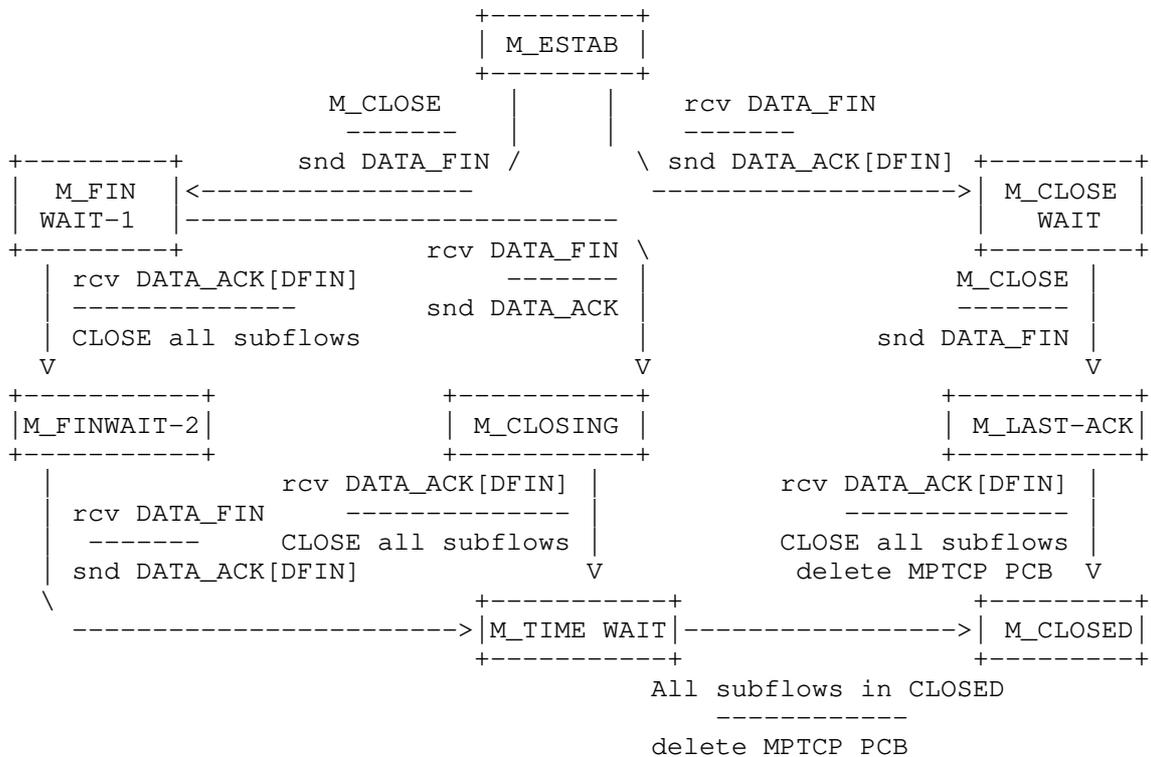


Figure 22: Finite State Machine for Connection Closure

Appendix E. Changes from RFC6824

This section lists the key technical changes between RFC6824, specifying MPTCP v0, and this document, which obsoletes RFC6824 and specifies MPTCP v1. Note that this specification is not backwards compatible with RFC6824.

- o The document incorporates lessons learnt from the various implementations, deployments and experiments gathered in the documents "Use Cases and Operational Experience with Multipath TCP" [RFC8041] and the IETF Journal article "Multipath TCP Deployments" [deployments].
- o Connection initiation, through the exchange of the MP_CAPABLE MPTCP option, is different from RFC6824. The SYN no longer includes the initiator's key, allowing the MP_CAPABLE option on the SYN to be shorter in length, and to avoid duplicating the sending of keying material.

- o This also ensures reliable delivery of the key on the MP_CAPABLE option by allowing its transmission to be combined with data and thus using TCP's in-built reliability mechanism. If the initiator does not immediately have data to send, the MP_CAPABLE option with the keys will be repeated on the first data packet. If the other end is first to send, then the presence of the DSS option implicitly confirms the receipt of the MP_CAPABLE.
- o In the Flags field of MP_CAPABLE, C is now assigned to mean that the sender of this option will not accept additional MPTCP subflows to the source address and port. This is an efficiency improvement, for example where the sender is behind a strict NAT.
- o In the Flags field of MP_CAPABLE, H now indicates the use of HMAC-SHA256 (rather than HMAC-SHA1).
- o Connection initiation also defines the procedure for version negotiation, for implementations that support both v0 (RFC6824) and v1 (this document).
- o The HMAC-SHA256 (rather than HMAC-SHA1) algorithm is used, as the algorithm provides better security. It is used to generate the token in the MP_JOIN and ADD_ADDR messages, and to set the initial data sequence number.
- o A new subflow-level option exists to signal reasons for sending a RST on a subflow (MP_TCP_RST Section 3.6), which can help an implementation decide whether to attempt later re-connection.
- o The MP_PRIO option (Section 3.3.8), which is used to signal a change of priority for a subflow, no longer includes the AddrID field. Its purpose was to allow the changed priority to be applied on a subflow other than the one it was sent on. However, it has been realised that this could be used by a man-in-the-middle to divert all traffic on to its own path, and MP_PRIO does not include a token or other security mechanism.
- o The ADD_ADDR option (Section 3.4.1), which is used to inform the other host about another potential address, is different in several ways. It now includes an HMAC of the added address, for enhanced security. In addition, reliability for the ADD_ADDR option has been added: the IPVer field is replaced with a flag field, and one flag is assigned (E) which is used as an 'Echo' so a host can indicate that it has received the option.
- o An additional way of performing a Fast Close is described, by sending a MP_FASTCLOSE option on a RST on all subflows. This

allows the host to tear down the subflows and the connection immediately.

- o In the IANA registry a new MPTCP subtype option, `MP_EXPERIMENTAL`, is reserved for private experiments. However, the document doesn't define how to use the subtype option.
- o A new Appendix discusses the usage of both the MPTCP and TCP Fast Open on the same packet (Appendix B).

Authors' Addresses

Alan Ford
Pexip

EEmail: alan.ford@gmail.com

Costin Raiciu
University Politehnica of Bucharest
Splaiul Independentei 313
Bucharest
Romania

EEmail: costin.raiciu@cs.pub.ro

Mark Handley
University College London
Gower Street
London WC1E 6BT
UK

EEmail: m.handley@cs.ucl.ac.uk

Olivier Bonaventure
Universite catholique de Louvain
Pl. Ste Barbe, 2
Louvain-la-Neuve 1348
Belgium

EEmail: olivier.bonaventure@uclouvain.be

Christoph Paasch
Apple, Inc.
Cupertino
US

EMail: cpaasch@apple.com

MPTCP Working Group
Internet-Draft
Intended status: Experimental
Expires: April 16, 2016

C. Paasch
A. Biswas
D. Haas
Apple, Inc.
October 14, 2015

Making Multipath TCP robust for stateless webrowsers
draft-paasch-mptcp-syncookies-02

Abstract

This document proposes a modification of the MPTCP handshake that allows it to work efficiently with stateless servers. We first identify the issues around stateless connection establishment using SYN-cookies. Further, we suggest an extension to Multipath TCP to overcome these issues and discuss alternatives.

As a side-effect, the proposed modification to the handshake opens the door to reduce the size of the MP_CAPABLE option in the SYN. This reduces the growing pressure on the TCP-option space in the SYN-segment, giving space for future extensions to 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 April 16, 2016.

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

1. Introduction	2
2. Problem statement	3
3. Proposal	4
3.1. Loss of the third ACK	4
3.1.1. MP_CAPABLE_ACK specification	5
3.1.2. TCP Fast Open	8
3.1.3. Negotiation	8
3.1.4. DATA_FIN	8
3.1.5. Server sending data	8
3.1.6. Middlebox considerations	9
3.2. Loss of the first data segment	10
4. Alternative solutions	11
5. IANA Considerations	11
6. Security Considerations	11
7. Acknowledgments	12
8. References	12
8.1. Normative References	12
8.2. Informative References	12
Authors' Addresses	13

1. Introduction

During the establishment of a TCP connection, a server must create state upon the reception of the SYN [RFC0793]. Specifically, it needs to generate an initial sequence number, and reply to the options indicated in the SYN. The server typically maintains in-memory state for the embryonic connection, including state about what options were negotiated, such as window scale factor [RFC7323] and the maximum segment size. It also maintains state about whether SACK [RFC2018] and TCP Timestamps were negotiated during the 3-way handshake.

Attackers exploit this state creation on the server through the SYN-flooding attack. Indeed, an attacker only needs to emit SYN segments with different 4-tuples (source and destination IP addresses and port numbers) in order to make the server create the state and thus consume its memory, while the attacker itself does not need to maintain any state for such an attack [RFC4987].

A common mitigation of this attack is to use a mechanism called SYN-cookies. SYN-cookies rely on the fact that a TCP-connection echoes back certain information that the server puts in the SYN/ACK during the three-way handshake. Notably, the sequence-number is echoed back in the acknowledgment field as well as the TCP timestamp value inside the timestamp option. When generating the SYN/ACK, the server generates these fields in a verifiable fashion. Typically, servers use the 4-tuple, the client's sequence number plus a local secret (which changes over time) to generate the initial sequence number by applying a hashing function to the aforementioned fields. Further, setting certain bits either in the sequence number or the TCP timestamp value allows to encode for example whether SACK has been negotiated and what window-scaling has been received [M08]. Upon the reception of the third ACK, the server can thus verify whether the acknowledgment number is indeed the reply to a SYN/ACK it has generated (using the 4-tuple and the local secret). Further, it can decode from the timestamp echo reply the required information concerning SACK, window scaling and MSS-size.

In case the third ACK is lost during the 3-way handshake of TCP, stateless servers only work if it's the client who initiates the communication by sending data to the server - which is commonly the case in today's application-layer protocols. As the data segment includes the acknowledgement number for the original SYN/ACK as well as the TCP timestamp value, the server is able to reconstruct the connection state even if the third ACK is lost in the network. If the very first data segment is also lost, then the server is unable to reconstruct the connection state and will respond to subsequent data sent by the client with a TCP Reset.

Multipath TCP (MPTCP [RFC6824]) is unable to reconstruct the MPTCP level connection state if the third ack is lost in the network (as explained in the following section). If the first data segment from the client reaches the server, the server can reconstruct the TCP state but not the MPTCP state. Such a server can fallback to regular TCP upon the loss of the third ACK. MPTCP is also prone to the same problem as regular TCP if the first data segment is also lost.

In the following section a more detailed assessment of the issues with MPTCP and TCP SYN-cookies is presented. Section 3 then shows how these issues might get solved.

2. Problem statement

Multipath TCP adds additional state to the 3-way handshake. Notably, the keys must be stored in the state so that later on new subflows can be established as well as the initial data sequence number is known to both hosts. In order to support stateless servers,

Multipath TCP echoes the keys in the third ACK. A stateless server thus can generate its own key in a verifiable fashion (similar to the initial sequence number), and is able to learn the client's key through the echo in the third ACK. The generation of the key is implementation-specific. An example of such a key-generation would be: `Key_Server = Hash(5-tuple, server's subflow sequence number, local_secret)`. The reliance on the third ACK however implies that if this segment gets lost, then the server cannot reconstruct the state associated to the MPTCP connection. Indeed, a Multipath TCP connection is forced to fallback to regular TCP in case the third ACK gets lost or has been reordered with the first data segment of the client, because it cannot infer the client's key from the connection and thus won't be able to generate a valid HMAC to establish new subflows nor does it know the initial data sequence number. In the remainder of this document we refer to the aforementioned issue as "Loss of the third ACK".

Stateless servers also are unable to recover connection state when the third ack and the first data segment are lost. This issue, outlined hereafter, happens even when regular TCP is being used. In case the client is sending multiple segments when initiating the connection, it might be that the third ack as well as the first data segment get lost. Thus, the server only receives the second data segment and will try to reconstruct the state based on this segment's 4-tuple, sequence number and timestamp value. However, as this segment's sequence number has already gone beyond the client's initial sequence number, it will not be able to regenerate the appropriate SYN-cookie and thus the verification will fail. The server effectively cannot infer that the sequence number in the segment has gone beyond TCP's initial sequence number. This will make the server send a TCP reset as it appears to the server that it received a segment for which no SYN cookie was ever generated.

3. Proposal

This section shows how the above problems might be solved in Multipath TCP.

3.1. Loss of the third ACK

In order to make Multipath TCP robust against the loss of the third ACK when SYN-cookies are being deployed on servers, we must make sure that the state-information relevant to Multipath TCP reaches the server in a reliable way. If the client is initiating the data transfer to the server (this data is being delivered reliably through TCP) the state-information could be delivered together with this data and thus is implicitly reliably sent to the server - when the data reaches the server, the state-information reaches the server as well.

We achieve this by adding another variant to the MP_CAPABLE option, differentiated by the length of it (we call this option MP_CAPABLE_ACK in the remainder of this document). It is solely sent on the very first data segment from the client to the server. This option serves the dual purpose of conveying the client's and server's key as well as the DSS mapping which would otherwise have been sent in a DSS option on the first data segment.

Making the MP_CAPABLE in the third ACK reliable opens the door for another improvement in MPTCP. In fact, the client doesn't need to send its own key in the SYN anymore (it will send it reliably in the MP_CAPABLE_ACK). Thus, the MP_CAPABLE option in the SYN segment can avoid adding the key, reducing the option-space requirement of the MP_CAPABLE down to 4 bytes. This is a major improvement as the option-space in the SYN segment is very limited, and allows a TCP connection to negotiate future extensions in the SYN.

As this change is a major extension to Multipath TCP, we require that the version number of the MP_CAPABLE is increased. Further details on the negotiation are presented in Section 3.1.3. The following is a detailed description of the option format and the suggested handshake.

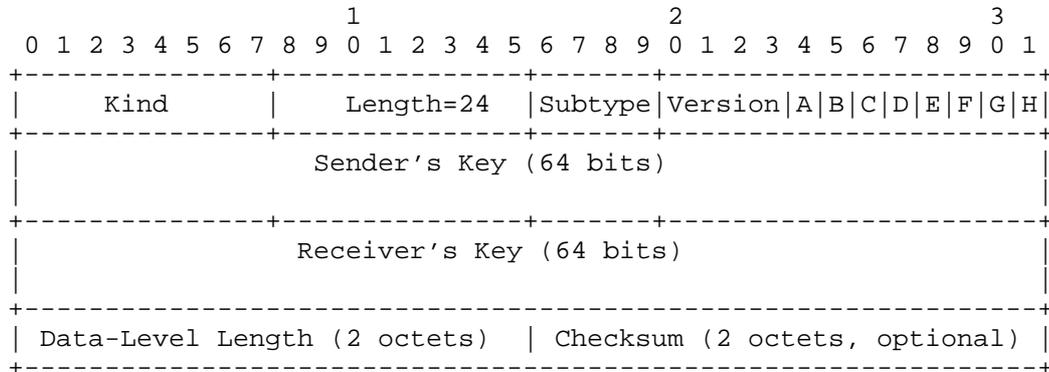
3.1.1.1. MP_CAPABLE_ACK specification

We suggest to remove the key from the MP_CAPABLE in the SYN-segment. The format of the MP_CAPABLE remains the same (with the bits A to H as well as the version number), with the difference that the key is no more present. Hosts are able to differentiate between the different MP_CAPABLE options through the length-field of the TCP-option.

The MP_CAPABLE option in the SYN/ACK as well as the third ACK (which does not contain any data) remain unmodified from RFC6824.

The MP_CAPABLE_ACK option (shown in Figure 1) contains the same set of bits A to H as well as the version number, like the MP_CAPABLE option. Further, the option includes the data-level length as well as the checksum (in case it has been negotiated during the 3-way handshake). This allows the server to reconstruct the mapping and deliver the data to the application. It must be noted that the information inside the MP_CAPABLE_ACK is less explicit than a DSS option. Notably, the data-sequence number, data acknowledgment as well as the relative subflow-sequence number are not part of the MP_CAPABLE_ACK. Nevertheless, the server is able to reconstruct the mapping because the MP_CAPABLE_ACK is guaranteed to only be sent on the very first data segment. Thus, implicitly the relative subflow-

sequence number equals 1 as well as the data-sequence number, which is equal to the initial data-sequence number.

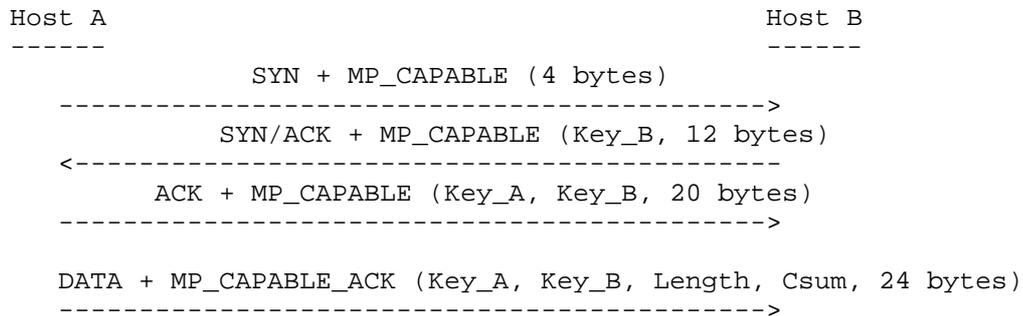


Format of the MP_CAPABLE_ACK option.

Figure 1

The handshake (depicted in Figure 2) starts with the client sending the MP_CAPABLE option to the server inside the SYN. The client is not required to having generated its key already at this point.

Upon reception of this SYN-segment, a stateful server generates a random key and replies with a SYN/ACK. If the server behaves in a stateless manner it has to generate it's own key in a verifiable fashion. This verifiable way of generating the key can be done by using a hash of the 4-tuple, sequence number and a local secret (similar to what is done for the TCP-sequence number [RFC4987]). It will thus be able to verify whether it is indeed the originator of the key echoed back in the MP_CAPABLE_ACK option. When generating this verifiable key, the server has to ensure that the token derived from this key is locally unique (Section 3.1 of RFC6824). If there is already an MPTCP-session with such a token, the server must fallback to regular TCP by not sending an MP_CAPABLE in the third ACK.



The modified MPTCP-handshake only consumes 4 bytes in the SYN.

Figure 2

To complete the three-way handshake, the client has to reply with a third ACK and the MP_CAPABLE option (with both keys as defined in RFC6824). If the client has already data to send, it can even avoid sending the third ACK, and immediately send the data together with the MP_CAPABLE_ACK. Otherwise, the client sends the MP_CAPABLE_ACK as soon as the application writes data on the socket.

The goal of the third ACK (with the MP_CAPABLE) as well as the MP_CAPABLE_ACK is to convey the client's key to the server. An indication for the client that the server received the key is when the server issues a DATA_ACK inside the DSS-option (even if this DATA_ACK does not acknowledge any data). Thus, as long as the client has not sent an MP_CAPABLE_ACK with data, it must add the MP_CAPABLE option in each (non-data) segment sent to the server. It must do this until it either did send an MP_CAPABLE_ACK or until it received a DATA_ACK from the server. The reason for this is explained in Section 3.1.2 and Section 3.1.5. Combining the MP_CAPABLE with the DATA_ACK will require 20 + 8 bytes, which still leaves 12 bytes for the TCP timestamp option.

Finally, the server must send a duplicate acknowledgment to the client upon reception of the client's key. This, to convey to the client that it successfully received the MP_CAPABLE(_ACK) option. It must be noted that this 4-way handshake does not prevent the client to send data before the reception of this fourth acknowledgment.

This mechanism of sending the MP_CAPABLE with a DATA_ACK until the server acknowledges it, introduces additional complexity to the handshake. However, we consider the gain of reducing the MP_CAPABLE option in the SYN-segment as significant enough, that it is worth to accept this added complexity.

3.1.2. TCP Fast Open

If TCP Fast Open [RFC7413] is being used in combination with Multipath TCP [I-D.barre-mptcp-tfo], the server is allowed to send data right after the SYN/ACK, without the need to wait for the third ACK. The server sending this data cannot include a DATA_ACK option inside the DSS option as it does not yet know the client's key. This is not an issue as the DATA_ACK is optional in the DSS option.

However, the client receiving this data sent by the server will have to acknowledge it with a DATA_ACK. As specified above, the client must also send an MP_CAPABLE option on this acknowledgment as it didn't yet receive a DATA_ACK from the server.

3.1.3. Negotiation

We require a way for the hosts to negotiate support for the suggested handshake. As we modify the size of the MP_CAPABLE, our proposal relies on a new version of MPTCP. The client requests this new version of MPTCP during the MP_CAPABLE exchange (it remains to be defined by the IETF which version of MPTCP includes the MP_CAPABLE_ACK option). If the server supports this version, it replies with a SYN/ACK including the MP_CAPABLE and indicating this same version.

3.1.4. DATA_FIN

As the MP_CAPABLE_ACK option includes the same bitfields as the regular MP_CAPABLE, there is no space to indicate a DATA_FIN as is done in the DSS option. This implies that a client cannot send a DATA_FIN together with the first segment of data. Thus, if the server requests the usage of MP_CAPABLE_ACK through the C-bit, the client must send a separate segment with the DSS-option, setting the DATA_FIN-flag to 1, after it has sent the data-segment that includes the MP_CAPABLE_ACK option.

3.1.5. Server sending data

The MP_CAPABLE_ACK version can only be sent by the client if it actually has data to send. One question that this raises is how the proposal will work when the server is the first one to send data to the client. In the following we describe how the handshake will still work when servers behave in a stateless and stateful manner.

For stateless servers the same issue arises as well for regular TCP. Upon loss of the third ACK, the server cannot complete the three-way handshake. Thus, stateless servers that begin the application level protocol by emitting data rely on the fact that the third ACK is

received (irregardless of whether MPTCP is used or not). Thus, this implies that the server also will receive the MP_CAPABLE with this third ACK.

Stateful servers will retransmit the SYN/ACK until the third ACK (including the MP_CAPABLE) has been received. This will thus provide to the server the client's key. When the client eventually sends its own first data segment to the server, it actually does not have to use the MP_CAPABLE_ACK option because the server already did send a DATA_ACK to the client.

3.1.6. Middlebox considerations

Multipath TCP has been designed with middleboxes in mind and so the MP_CAPABLE_ACK option must also be able to go through middleboxes. The following middlebox behaviors have been considered and MP_CAPABLE_ACK acts accordingly across these middleboxes:

- o Removing MP_CAPABLE_ACK-option: If a middlebox strips the MP_CAPABLE_ACK option out of the data segment, the server receives data without a corresponding mapping. As defined in Section 3.6 of [RFC6824], the server must then do a seamless fallback to regular TCP.
- o Coalescing segments: A middlebox might coalesce the first and second data segment into one single segment. While doing so, it might remove one of the options (either MP_CAPABLE_ACK or the DSS-option of the second segment because of the limited 40 bytes TCP option space). There are two cases to consider:
 - * If the DSS-option is not included in the segment, the second half of the payload is not covered by a mapping. Thus, the server will do a seamless fallback to regular TCP as defined by [RFC6824] in Section 3.6. This fallback will trigger because RFC6824 specifies that during the beginning of a connection (as long as the path has not been proven to let the MPTCP-options unmodified in both directions) a seamless fallback to regular TCP must be done by stopping to send DATA_ACKs to the client.
 - * If the MP_CAPABLE_ACK option is not present, then the DSS-option provides an offset of the TCP sequence number. As the server behaves statelessly it can only assume that the present mapping belongs to the first byte of the payload (similar to what is explained in detail in Section 3.2). As this however is not true, it will calculate an incorrect initial TCP sequence number and thus reply with a TCP-reset as the SYN-cookie is invalid. As such kind of middleboxes are very rare we consider this behavior as acceptable.

- o Splitting segments: A TCP segmentation offload engine (TSO) might split the first segment in smaller segments and copy the MP_CAPABLE_ACK option on each of these segments. Thanks to the data-length value included in the MP_CAPABLE_ACK option, the server is able to detect this and correctly reconstructs the mapping. In case the first of these splitted segments gets lost, the server finds itself in a situation similar to the one described in Section 2. The TCP sequence number doesn't allow anymore to verify the SYN-cookie and thus a TCP reset is sent. This behavior is the same as for regular TCP.
- o Payload modifying middlebox: In case the middlebox modifies the payload, the DSS-checksum included in the MP_CAPABLE_ACK option allows to detect this and will trigger a fallback to regular TCP as defined in [RFC6824].

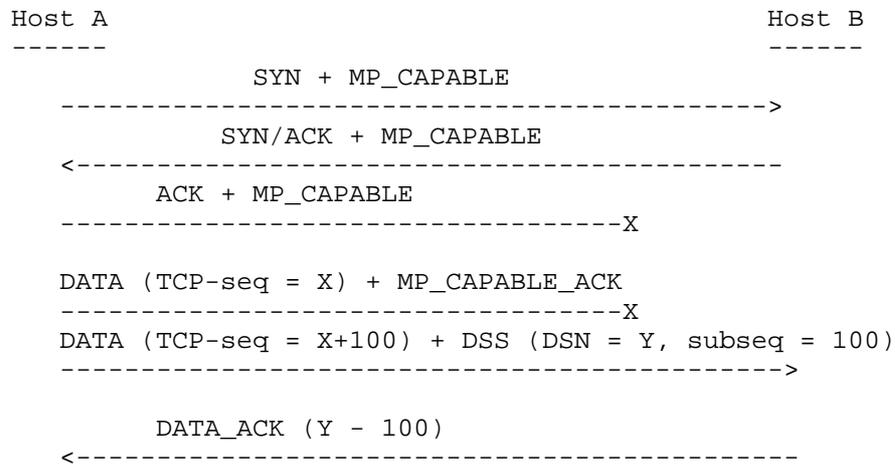
3.2. Loss of the first data segment

Section 2 described the issue of losing the first data segment of a connection while TCP SYN-cookies are in use. The following outlines how Multipath TCP actually allows to fix this particular issue.

Consider the packet-flow of Figure 3. Upon reception of the second data segment, the included data sequence mapping allows the server to actually detect that this is not the first segment of a TCP connection. Indeed, the relative subflow sequence number inside the DSS-mapping is actually 100, indicating that this segment is already further ahead in the TCP stream. This allows the server to actually reconstruct the initial sequence number based on the sequence number in the TCP-header $((X+100) - 100)$ that has been provided by the client and verify whether its SYN-cookie is correct. Thus, no TCP-reset is being sent - in contrast to regular TCP, where the server cannot verify the SYN-cookie. The server knows that the received segment is not the first one of the data stream and thus it can store it temporarily in the out-of-order queue of the connection. It must be noted that the server is not yet able to fully reconstruct the MPTCP state. In order to do this it still must await the MP_CAPABLE_ACK option that is provided in the first data segment.

The server responds to the out-of-order data with a duplicate ACK. The duplicate ACK may also have SACK data if SACK was negotiated. However, if this duplicate ACK does not have an MPTCP level Data ACK, the client may interpret this as a fallback to TCP. This is because the client cannot determine if an option stripping middlebox removed the MPTCP option on TCP segments after connection establishment. So even though the server has not fully recreated the MPTCP state at this point, it should respond with a Data ACK set to the Data Sequence Number $Y-100$. The client's TCP implementation may

retransmit the first data segment after a TCP retransmit timeout or it may do so as part of an Early Retransmit that can be triggered by an ACK arriving from the server.



Multipath TCP's DSS option allows to handle the loss of the first data segment as the host can infer the initial sequence number.

Figure 3

4. Alternative solutions

An alternative solution to creating the MP_CAPABLE_ACK option would have been to emit the MP_CAPABLE-option together with the DSS-option on the first data segment. However, as the MP_CAPABLE option is 20 bytes long and the DSS-option (using 4-byte sequence numbers) consumes 16 bytes, a total of 36 bytes of the TCP option space would be consumed by this approach. This option has been dismissed as it would prevent any other TCP option in the first data segment, a constraint that would severely limit TCP's extensibility in the future.

5. IANA Considerations

Our proposal requires the change of the MPTCP-version number.

6. Security Considerations

Sending the keys in a reliable way after the three-way handshake implies that there is a larger window during which an on-path attacker might modify the keys that are being sent in the MP_CAPABLE_ACK. However, we do not think that this can actually be

considered as a security issue. If an attacker modifies the keys, the outcome will be that the client and the server won't agree anymore on the data-sequence numbers. The data-flow will thus stall. Considering that the attacker has to be an active on-path attacker to launch this attack, he has already other means of interfering with the connection. Thus, this attack is considered as irrelevant.

Further, if servers implement the proposal from Section 3.2, to handle the scenario where the first data-segment is lost, the incoming segments need to be stored in the out-of-order queue. The server will store these segments without having verified the key that the client provides in the MP_CAPABLE option. This might be considered as a security risk where an attacker could consume buffer space in the server. It must be noted however that in order to achieve this, the attacker needs to correctly guess the SYN-cookie so that the verification described in Section 3.2 is successful. As MPTCP does not try to be more secure than regular TCP, this thread can be considered acceptable, as it uses the same level of security as regular TCP's SYN-cookies. Nevertheless, servers are free to avoid storing those segments in the out-of-order queue if the thread is considered important enough.

7. Acknowledgments

We would like to thank Olivier Bonaventure, Yoshifumi Nishida and Alan Ford for their comments and suggestions on this draft.

8. References

8.1. Normative References

- [RFC4987] Eddy, W., "TCP SYN Flooding Attacks and Common Mitigations", RFC 4987, August 2007.
- [RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013.

8.2. Informative References

- [I-D.barre-mptcp-tfo] Barre, S., Detal, G., and O. Bonaventure, "TFO support for Multipath TCP", draft-barre-mptcp-tfo-01 (work in progress), January 2015.
- [M08] McManus, P., "Improving syncookies", 2008, <<http://lwn.net/Articles/277146/>>.

- [RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC2018] Mathis, M., Mahdavi, J., Floyd, S., and A. Romanow, "TCP Selective Acknowledgment Options", RFC 2018, October 1996.
- [RFC7323] Borman, D., Braden, B., Jacobson, V., and R. Scheffenegger, "TCP Extensions for High Performance", RFC 7323, September 2014.
- [RFC7413] Cheng, Y., Chu, J., Radhakrishnan, S., and A. Jain, "TCP Fast Open", RFC 7413, December 2014.

Authors' Addresses

Christoph Paasch
Apple, Inc.
Cupertino
US

Email: cpaasch@apple.com

Anumita Biswas
Apple, Inc.
Cupertino
US

Email: anumita_biswas@apple.com

Darren Haas
Apple, Inc.
Cupertino
US

Email: dhaas@apple.com

MPTCP working Group
Internet Draft
Intended status: Informational
Expires: Dec 26, 2015

A. Palanivelan
Verizon Labs
H. Chetan
Verizon Labs
S. Senthil
Cisco Systems
June 29, 2015

MPTCP Implementation (Challenges) Usecases
draft-palanivelanchetan-mptcp-challenges-usecases-00

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on December 26, 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.

Abstract

MPTCP Intends to address a wide range of issues, with minimal implementation tweaks. Though this works in a range of use cases, there are some use cases, where some standard implementation recommendations could help. The Purpose of this draft is to document the use cases, where there are opportunities for standard implementation recommendations.

Table of Contents

- 1. Introduction.....3
- 2. MPTCP Implementation Challenges -Usecases.....3
 - 2.1. Impact on existing monitoring and security tools.....3
 - 2.2. Fragmentation.....3
 - 2.3. Aggregation, Persistence & Resilience.....3
 - 2.4. Address Hopping.....4
 - 2.5. Security Issues.....4
 - 2.6. MPTCP for DATA CENTRES.....4
 - 2.7. MPTCP performance in presence and absence of proxies.....4
- 3. Security Considerations.....4
- 4. IANA Considerations.....4
- 5. Conclusions.....4
- 6. References.....5
 - 6.1. Normative References.....5
 - 6.2. Informative References.....5
- 7. Acknowledgments.....5

1. Introduction

There are several existing MPTCP Implementations and no doubt there is a growing demand to adopt this in addressing a wide range of issues.

This draft documents use cases where there are implementation challenges and would need standard recommendations. The Opportunities (design recommendations) to address one or set of use cases, independently is not in the scope of this document.

2. MPTCP Implementation Challenges -Usecases

2.1. Impact on existing monitoring and security tools

If an endpoint sending data wants to add paths to an MPTCP session, it notifies the receiving server, which then opens up the new path. That's a contradiction -- it's meant to be a stream sending data (envision that as data being sent from "left" to "right" on a diagram), but the connection is originating in the other direction, right to left. From the point of view of a network tool, that new stream is "outbound, but incoming," That could cause confusion in some monitoring and security tools.

2.2. Fragmentation

In MPTCP since the party can choose how to divide up data it can do fragmentation attacks very easily, for instance by sending every alternate byte down alternate paths vi different interfaces. Because of technical details of how the mapping occurs, this wouldn't be very efficient, however if both ends are doing this then there may be no single point on the network which can observe the traffic from both flows. This is the unique danger of a multipath fragmentation attack reassembly capability alone is not sufficient to observe all traffic.

2.3. Aggregation, Persistence & Resilience

In the immediate term however, MPTCP is designed to provide reliable delivery over multiple paths and to detect where interference happens and avoid that path. This may cause issues for organizations who use active or transparent security measures. If a client can find a path that supports MPTCP it will prefer it over the recommended path, leading to both avoidance of security measures and the potential for new forms of phishing attacks.

2.4. Address Hopping

An endpoint can add and remove addresses at will throughout the connection, with no requirement that any connection remain available, just that the chain of connections remains so. This adds additional complexity for monitoring and state handling, for now, we need to watch every connection, every possible connection, and track every advertised address.

2.5. Security Issues

Many organizations use network security which is highly dependent on tight perimeters, flow monitoring and data inspection. Where network security technology is not ready, it may be significantly less effective against MPTCP traffic

2.6. MPTCP for DATA CENTRES

2.7. MPTCP performance in presence and absence of proxies

3. Security Considerations

None

4. IANA Considerations

None

5. Conclusions

None

6. References

6.1. Normative References

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

6.2. Informative References

[RFC6182] Ford, A., Raiciu, C., Handley, M., Barre, S., and J.Iyengar, "Architectural Guidelines for Multipath TCP Development", RFC 6182, March 2011.

[RFC6356] Raiciu, C., Handley, M., and D. Wischik, "Coupled Congestion Control for Multipath Transport Protocols", RFC6356, October 2011.

7. Acknowledgments

We like to appreciate and thank all contributions to this document, by means of useful feedback and inputs

Authors' Addresses

Chetan Harsha
R&D Software Engineer, Verizon Labs
Bangalore, India

Email: chetan.harsha@verizon.com

Senthil sivakumar
Principal Engineer, Cisco systems
Durham, NC

Email: ssenthil@cisco.com

Palanivelan Appanasamy
DMTS-Engineering, Verizon Labs
Bangalore, India

Email: palanivelan.appanasamy@verizon.com

INTERNET-DRAFT
Intended Status: Standards Track
Expires: January 2, 2016

X.We
C.Xiong
Huawei Technologies
E. Lopez
Fortinet
July 1, 2015

MPTCP proxy mechanisms
draft-wei-mptcp-proxy-mechanism-02

Abstract

Multipath TCP provides the ability to simultaneously use multiple paths between peers for a TCP/IP session, and it could improve resource usage within the network and, thus, improve user experience through higher throughput and improved resilience to network failure. This document discusses the mechanism of a new network entity, named MPTCP proxy, which is aimed to assist MPTCP capable peer to use MPTCP session in case of one of the peers not being MPTCP capable or to act as an aggregation point for subflows.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Copyright and License 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

1	Introduction	3
2	Terminology	3
3	MPTCP Proxy models	4
4	MPTCP Proxy Solutions	5
4.1	Mechanisms for on-path MPTCP proxy	5
4.2	Mechanisms for off-path MPTCP proxy	7
5	Conclusion	9
6	Security Considerations	10
7	IANA Considerations	10
8	References	10
8.1	Normative References	10
8.2	Informative References	11
	Authors' Addresses	11

1 Introduction

Nowadays, the volume of mobile devices, e.g. smart phone, has increased greatly, and most of these devices have more than one interface for network communication, for example it's very common for a smart phone to have a cellular network interface and a WLAN interface; at the same time, multi-homing scenarios have been more and more common. All these situations provide a good pre-condition for the implementation of MPTCP [MPTCP Protocol]. Some network operators also show interests in MPTCP, they want to utilize MPTCP's multipath feature to realize optimization of their network performances, such as resource pooling, network mobility etc.

But there are still some barriers existing for the promotion of MPTCP, and one of them is that now almost all of the ICP (Internet Content Provider) servers on the Internet are traditional TCP servers and there seems no motivation for these traditional servers to embed MPTCP into their protocol stack, this situation leads to the fact that when communicating with these servers the MPTCP capable devices have to fall back to traditional TCP and cannot fully utilize their MPTCP capability.

Besides, the multipath feature of MPTCP protocol brings impacts on the performances of some kinds of network middleboxes which are deployed to enhance network performance or to provide traffic optimization for network traffic. For example, middleboxes, such as HTTP proxy, video/audio optimizer and firewall are deployed enroute by network operators to provide performance enhancements, and all of these middleboxes need to have knowledge about the entire content of the traffic flow in order to function properly on the flow. But for MPTCP traffic, it is likely that only a part of subflows traverse the middlebox, and leads these middleboxes to be blind about the traffic, and the result would be that the endhost could not benefit from performance enhancement service or the traffic from endhost could be blocked by firewall because the firewall cannot trust the traffic. A more detailed description of MPTCP's impacts on middleboxes can be found in [Lopez]. For all the middlebox scenarios, we can conclude a basic requirement that the MPTCP traffic should be able to aggregate at middlebox.

To support the use of MPTCP session between a MPTCP host and a TCP host, and to make MPTCP traffic get benefits from network middlebox that providing performance enhancement, this document defines a new entity named MPTCP proxy (or proxy for abbreviation).

2 Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",

"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

MPTCP proxy (proxy): An entity used to support MPTCP session between MPTCP capable host and non-MPTCP capable host.

ICP: Internet Content Provider.

3 MPTCP Proxy models

To support the use of MPTCP session between MPTCP host and TCP host, or to help to aggregate MPTCP subflows, there are mainly two models of proxy for different scenarios: the first one is that the proxy is deployed on the common direct routing path of traffic from different access network, and this kind of proxy is referred as on-path proxy, an example is shown in Figure 1; the second one is that the proxy locates only on the direct routing path of traffic from one of the access networks the MPTCP capable host attached to, and this kind of proxy is referred as off-path proxy, an example is shown in Figure 2.

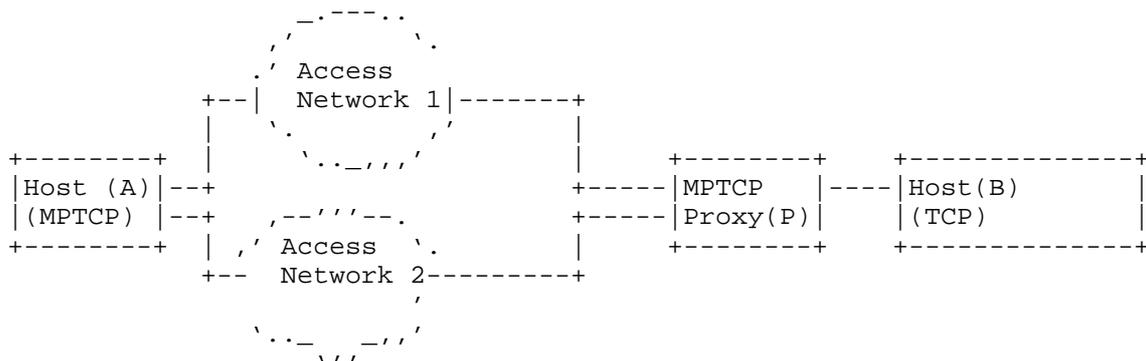


Figure 1: Scenario of on-path proxy deployment

For the scenario shown in Figure 1, the MPTCP capable host A has two network interfaces and connects to two access networks simultaneously through the two interfaces. In this case, the proxy is located on the path shared by the two access networks' traffic, for example, the proxy could be deployed by Host B (e.g. OTT server) side.

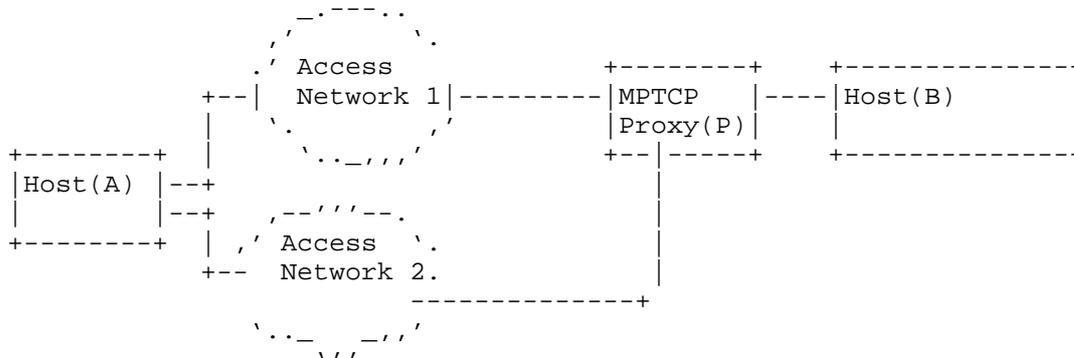


Figure 2: Scenario of off-path proxy deployment

For the scenario shown in Figure 2, MPTCP proxy is located only on the direct routing path of traffic from one of the access networks the MPTCP capable host attached to, for example, the proxy could be located at aggregation point such as firewall. As shown in Figure 2, the MPTCP proxy is located on the natural routing path of traffic from access network 1, but not on the natural routing path of traffic from access network 2, which means when host A communicates using MPTCP with host B, the subflow through access network 2 will not be naturally routed to MPTCP proxy.

The MPTCP communication in this scenario could occur between a MPTCP host and a TCP host, or two MPTCP hosts.

The following sections will discuss the detailed mechanisms of on-path proxy and off-path proxy as introduced above.

4 MPTCP Proxy Solutions

4.1 Mechanisms for on-path MPTCP proxy

When the direct routing path of all the sub-flows of a MPTCP capable host pass through the same proxy, the proxy will act as on-path proxy, and the on-path proxy could be transparent to the end host, i.e. end host itself knows nothing about the existence of the proxy.

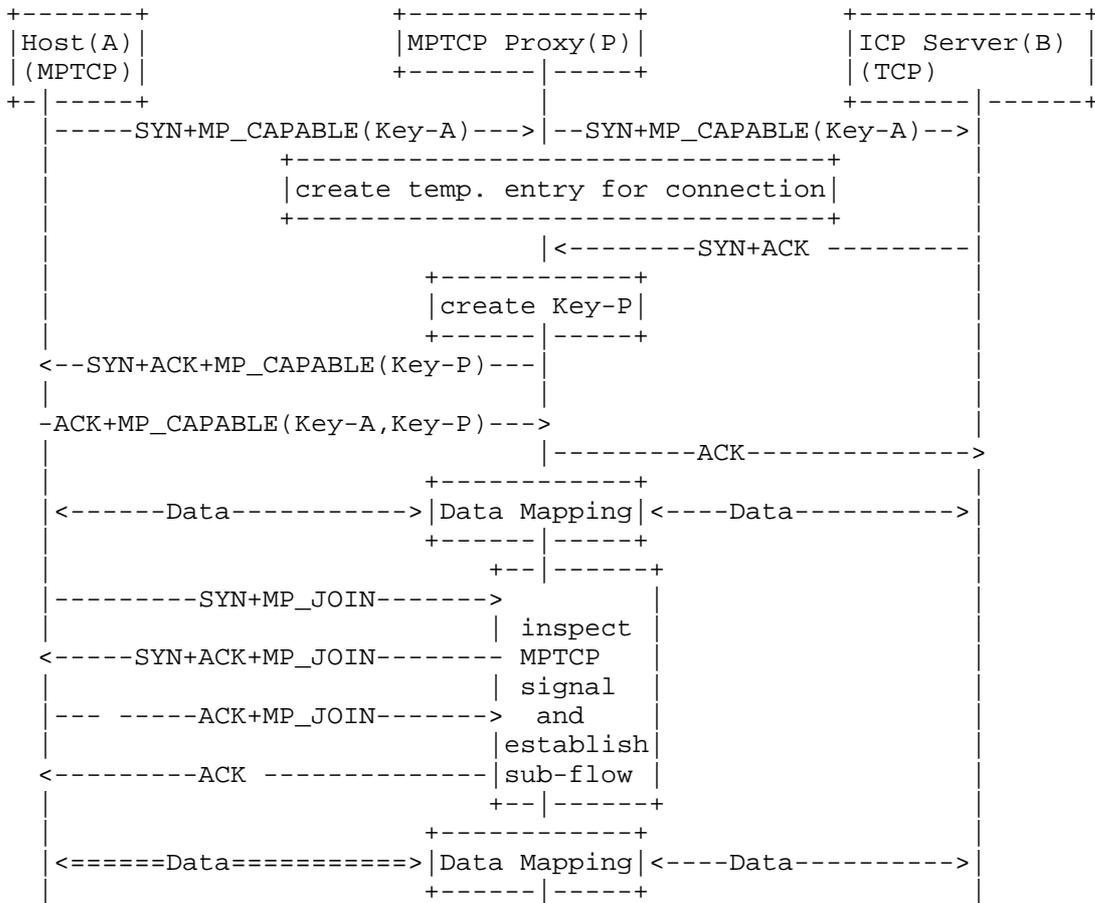


Figure 3: On-path proxy for connection between MPTCP Host and TCP Server

The function of on-path proxy could mainly be divided into three sub-functions: supporting for initial MPTCP capability negotiation, supporting for sub-flow establishment and data mapping. Figure 3 shows an example signal flow for on-path proxy. The following clauses focus on the description of each sub-function.

(1) Supporting for initial MPTCP capability negotiation

The MPTCP capable host starts a connection establishment procedure by sending the first handshake packet with MP_CAPABLE option, including Host's Key-A, to ICP server; proxy inspects the packet and creates a temporary entry, which will be used to match SYN/ACK response from ICP server, for the connection, then the proxy forwards the packet to ICP

server.

(2) Supporting for sub-flow establishment

After the initial MPTCP connection established, Host could choose to start a new MPTCP sub-flow. Because Host is unaware of the existence of proxy, so Host will start the new sub-flow with ICP server, i.e. the destination IP address of SYN/MP_JOIN packet is ICP server's IP address.

The proxy inspects sub-flow establishment signal packet, i.e. SYN/MP_JOIN, and decides whether it has provided proxy function for the MPTCP session through the token included in MP_JOIN. If proxy has provided proxy function for the MPTCP session, then it will provide proxy function for the sub-flow; otherwise proxy will not take any action on the establishment of sub-flow.

(3) Data mapping

Proxy implements two separate kinds of data mapping: forward mapping and reverse mapping. Forward mapping means mapping data from MPTCP session to TCP session; reverse mapping means mapping data from TCP session to MPTCP session. Figure 4 shows the data mapping function of proxy. In forward mapping, proxy maps data from all sub-flows belonging to MPTCP session to a single TCP flow in TCP session.

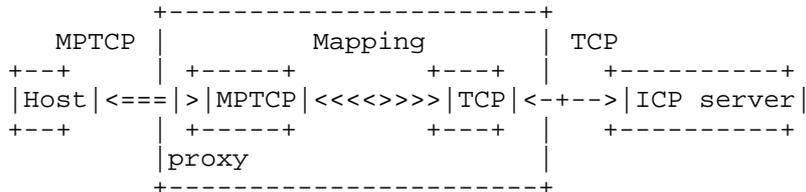


Figure 4: Data mapping function of proxy

4.2 Mechanisms for off-path MPTCP proxy

When proxy locates on the initial sub-flow's direct routing path, but some other sub-flow's direct routing path might not go through the same proxy, then proxy could act in off-path model. The main difference between on-path model proxy and off-path model proxy is that in off-path model proxy needs to steer sub-flows to proxy, and Host will start new sub-flow with proxy, but not with its peer host.

MP_CAPABLE option is included in SYN/ACK packet, then it means the ICP server is MPTCP capable and the proxy could choose whether to act as proxy or not for the connection, for example if the proxy wants to act as an aggregation point for MPTCP subflow traffics then it could choose to act proxy function for the MPTCP session; but if the purpose of proxy is just provide MPTCP support for communication between MPTCP host and legacy TCP host, then it could choose not to act proxy function; if no MP_CAPABLE option is included in SYN/ACK, the proxy will generate Key-P on behalf of ICP server to finish MPTCP connection with Host.

To avoid Host starts the establishment of sub-flow with ICP server's IP address, proxy notifies Host the existence of itself through sending a P flag in MP_CAPABLE option in SYN/ACK packet. When Host receives this P flag it SHOULD NOT start the new sub-flow with ICP server's IP address any more, but chooses to establish sub-flow with proxy after obtaining proxy's IP address.

There are reasons why a new P flag needs to be defined for explicit indication the existence of proxy, instead of implicitly inject the proxy into MPTCP session using existing MPTCP signaling, e.g. using ADD_ADDR/ADDR_JOIN to inform MPTCP host of proxy's IP address, and using REMOVE_ADDR to disable initial subflow between MPTCP host and its peer host: When a new subflow is to be established, the subflow management strategy should be considered. As stated in [MPTCP Experience], "The subflows are created immediately after the creation of the initial subflow", so MPTCP host might have started a new subflow before a REMOVE_ADDR is received, due to message delay or lost of REMOVE_ADDR, in that case the new subflow might be established between MPTCP host and its peer host but not between MPTCP host and proxy.

(2) Supporting for sub-flow establishment

In off-path model, after MPTCP capable Host has established the initial sub-flow in MPTCP session with the assistance of proxy, proxy could advertise its own IP address in ADD_ADDR option to Host, and then Host could establish new sub-flow with proxy.

(3) Data mapping

The data mapping function for off-path proxy is the same as the function described in on-path model.

5 Conclusion

This document provides two kinds of proxy modes, which could be used to support MPTCP capable Host in two different scenarios. For the first on-path MPTCP proxy, there is no need to modify the current MPTCP stack implementation of the host; for the off-path MPTCP proxy, it requires

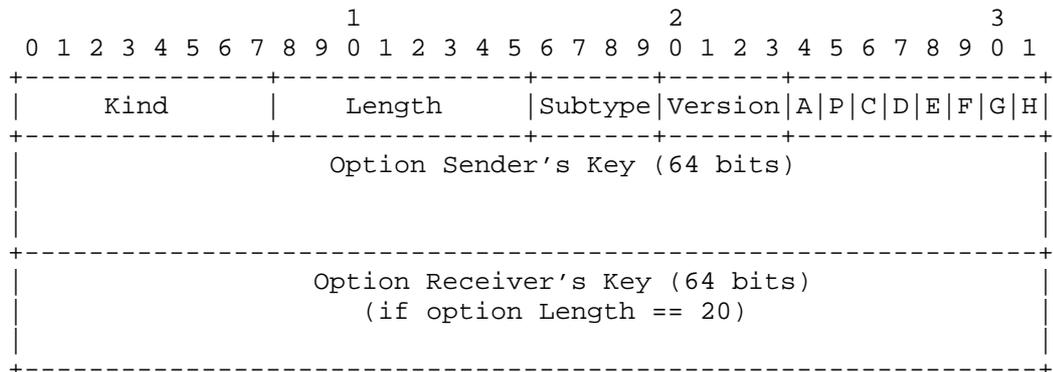
the MPTCP capable host needs to support a new defined P flag.

6 Security Considerations

The P flag provides a method for explicitly interpose a proxy function in MPTCP session, but this does not bring more security risks than MPTCP protocol itself, because even without P flag, an on-path middlebox could still interpose it in MPTCP session using existing MPTCP protocol signaling.

7 IANA Considerations

A new flag 'P' in MPTCP MP_CAPABLE option [MPTCP Protocol] needs to be defined refer to RFC 6824, Section 3.1. This flag is used by proxy to inform MPTCP capable host the existence of proxy, besides the 'P' flag could also be used to inform other potential MPTCP proxy its presence.



8 References

8.1 Normative References

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

8.2 Informative References

- [Deng] L.Deng, D.Liu, T.Sun. "draft-deng-mptcp-mobile-network-proxy-01", April 18, 2014
- [Lopez] E. Lopez. "draft-lopez-mptcp-middlebox-00", November 11, 2014
- [MPTCP Experience] O. Bonaventure et al. "draft-ietf-mptcp-experience-00". September 16, 2014

Authors' Addresses

Xinpeng Wei
Huawei Technoligies
Beijing, China
EMail: weixinpeng@huawei.com

Chunshan Xiong
Huawei Technoligies
Beijing, China
EMail: sam.xiongchunshan@huawei.com

Edward Lopez
Fortinet
899 Kifer Road
Sunnyvale, CA 94086
EMail: elopez@fortinet.com