

Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: October 22, 2012

R. Winter
University of Applied Sciences
Augsburg
A. Ripke
NEC Laboratories Europe
April 20, 2012

Multipath TCP Support for Single-homed End-systems
draft-wr-mptcp-single-homed-02

Abstract

Multipath TCP relies on the existence of multiple paths at the end-systems typically provided through different IP addresses obtained by different ISPs. While this scenario is certainly becoming increasingly a reality (e.g. mobile devices), currently most end-systems are single-homed (e.g. residential broadband customers). This memo describes mechanisms to make multiple paths available to multipath TCP via autoconfiguration that are not available at the end-systems but somewhere within the network.

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 October 22, 2012.

Copyright Notice

Copyright (c) 2012 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.	General Operation	4
3.	Other scenarios and extensions	5
4.	Alternative approaches	6
5.	Acknowledgements	6
6.	IANA Considerations	6
7.	Security Considerations	6
8.	References	7
8.1.	Normative References	7
8.2.	Informative References	7
	Authors' Addresses	7

1. Introduction

The IETF is specifying a multipath TCP (MPTCP) architecture and protocol where end-systems operate a modified standard TCP stack which allows packets of the same TCP connection to be sent via different paths to an MPTCP-capable destination ([[I-D.ietf-mptcp-multiaddressed](#)], [[RFC6182](#)]) where paths are defined by sets of source and destination IP addresses. Using multiple paths has a number of benefits such as an increased reliability of the transport connection and an effect known as resource pooling [[resource pooling](#)]. Most end-systems today do not have multiple paths/interfaces available in order to make use of multipath TCP, however further within the network multiple paths are the norm rather than the exception. This memo therefore describes a way how a network element that is multi-homed can somehow "extend" this multi-homing to end-systems, providing the multipath TCP benefits starting at that network element.

In order to illustrate the general mechanism we make use of a simple reference scenario shown in Figure 1.




Figure 1: Reference Scenario

The scenario in Figure 1 depicts e.g. a possible SOHO or enterprise setup where a gateway/router is connected to two ISPs and a DHCP server gives out leases to hosts connected to the local network. Note that both, the gateway and the DHCP server could be on the same device (similar to current home gateway implementations).

The host is running a multipath-capable IP stack, however it only has a single interface. The method described in the following sections is to let the host make use of the gateway's two interfaces without requiring modifications to the MPTCP implementation. In particular, we describe a way to autoconfigure the host.

2. General Operation

Multipath TCP distinguishes paths by their source and destination IP addresses. Assuming a certain level of path diversity in the Internet, using different source and destination IP addresses for a given subflow of a multipath TCP connection will, with a certain probability, result in different paths taken by packets of different subflows. Even in case subflows share a common bottleneck, the proposed multipath congestion control algorithm [RFC6356] will make sure that multipath TCP will play nicely with regular TCP flows.

In order to not require changes to the TCP implementation, we keep the above assumptions multipath TCP makes, i.e. working with different IP addresses to use different paths. Since the end-system is single-homed, all IP addresses are bound to the same physical interface. In our reference scenario in Figure 1, the host would receive more than one RFC1918 [RFC1918] private IP address from the DHCP server as depicted in Figure 2.




Figure 2: Gateway internals

The gateway that is shown in Figure 2 has received two IP addresses, one from each ISP that it is connected to. The NAT that the gateway is implementing needs to "map" each private IP address of the host consistently to a one of the addresses, each private IP to a different public IP. Packets sent by the host to the gateway are then routed based on the source address found in the packets as illustrated in the figure. In other words, depending on the source address of the host, the packets will either go through ISP 1 or ISP 2 and TCP will balance the traffic across those two links using its built-in congestion control mechanism.

The way the gateway has received its public IP addresses is not relevant. It could be via DHCP, IPCP or static configuration.

In order to configure the host via DHCP, we propose two new DHCP options. The first option "mp-avail" will be sent by single-homed multipath TCP-enabled clients in the "Parameter Request List". This will show the DHCP server that the client is multipath-capable. The DHCP server will answer with "mp-avail" and the option value is set to the number of additional interfaces the gateway can offer to the client (in our reference scenario that value would be 1; see Figure 3).




Figure 3: DHCP Sequence Diagram

Upon receipt of the "mp-avail" option from the server, the client can create up to n virtual interfaces, where n is the option value. Each virtual interface will contact the DHCP server and will include the "mp-range" option. The option value will tell the DHCP server that the client is requesting an IP out of an IP range that the gateway will be forwarding through a different interface.

The above has been implemented using the ISC DHCP server and client version 4.2.1 and the multipath TCP kernel patch 0.5 and a 2.6.36 Linux kernel.

3. Other scenarios and extensions

The reference scenario is only one conceivable setting. Other scenarios such as DSL broadband customers or mobile phones are

conceivable as well. As an example, take the DSL scenario. The home gateway could be provided with multiple IP addresses using extensions to IPCP. The home gateway in turn can then implement the DHCP server and gateway functionality as described before. More scenarios will be described in future versions of this document.

4. Alternative approaches

One alternative is that a DHCP server always sends n offers, where n is the number of interfaces at the gateway to different ISPs. The client could then accept all or a subset of these offers. This approach seems interesting in environments where there are multiple DHCP servers, one for each ISP connection (think multiple homegateways). However, accepting multiple offers based on a single DHCP request is not standard's compliant behavior. Also, to cater for a scenario that only contains a single DHCP server, server changes are needed in any case. Finally, correct routing is not always guaranteed in these scenarios.

An interesting alternative is the use of ECMP at the gateway for load distribution and let MPTCP use different port numbers for subflows. Assuming that ECMP is available at the gateway, this approach would work fine today. The only drawback of the approach is that it involves a little trial and error to find port numbers that actually hash to different paths used by ECMP.

5. Acknowledgements

Part of this work was supported by Trilogy (<http://www.trilogy-project.org>), a research project (ICT-216372) partially funded by the European Community under its Seventh Framework Program. The views expressed here are those of the author(s) only. The European Commission is not liable for any use that may be made of the information in this document.

6. IANA Considerations

Two new DHCP options are required by this version of this document.

7. Security Considerations

TBD.

8. References

8.1. Normative References

- [RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", [BCP 5](#), [RFC 1918](#), February 1996.

8.2. Informative References

- [I-D.ietf-mptcp-multiaddressed]
Ford, A., Raiciu, C., Handley, M., and O. Bonaventure,
"TCP Extensions for Multipath Operation with Multiple
Addresses", [draft-ietf-mptcp-multiaddressed-07](#) (work in
progress), March 2012.
- [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",
[RFC 6356](#), October 2011.
- [resource_pooling]
Wischik, D., Handley, M., and M. Bagnulo Braun, "The
Resource Pooling Principle", October 2008,
<<http://ccr.sigcomm.org/online/files/p47-handleyA4.pdf>>.

Authors' Addresses

Rolf Winter
University of Applied Sciences Augsburg
An der Hochschule 1
Augsburg 86161
Germany

Email: rolf.winter@hs-augsburg.de

Andreas Ripke
NEC Laboratories Europe
Kurfuersten-Anlage 36
Heidelberg 69115
Germany

Email: andreas.ripke@neclab.eu