Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 20, 2017

M. Boucadair
C. Jacquenet
Orange
T. Reddy
Cisco

April 18, 2017

DHCP Options for Network-Assisted Multipath TCP (MPTCP)
draft-boucadair-mptcp-dhc-07

Abstract

Because of the lack of Multipath TCP (MPTCP) support at the server side, some service providers now consider a network-assisted model that relies upon the activation of a dedicated function called MPTCP Conversion Point (MCP). Network-assisted MPTCP deployment models are designed to facilitate the adoption of MPTCP for the establishment of multi-path communications without making any assumption about the support of MPTCP by the communicating peers. MCPs located in the network are responsible for establishing multi-path communications on behalf of endpoints, thereby taking advantage of MPTCP capabilities to achieve different goals that include (but are not limited to) optimization of resource usage (e.g., bandwidth aggregation), of resiliency (e.g., primary/backup communication paths), and traffic offload management.

This document focuses on the explicit deployment scheme where the identity of the MPTCP Conversion Point(s) is explicitly configured on connected hosts. This document specifies DHCP (IPv4 and IPv6) options to configure hosts with Multipath TCP (MPTCP) parameters.

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].

Table of Contents

1. Introduction

One of the promising deployment scenarios for Multipath TCP (MPTCP, [RFC6824]) is to enable a Customer Premises Equipment (CPE) that is connected to multiple networks (e.g., DSL, LTE, WLAN) to optimize the usage of such resources. This deployment scenario relies on MPTCP Conversion Points (MCPs) located on both the CPE and network sides (Figure 1). The latter plays the role of traffic concentrator. An MCP terminates the MPTCP sessions established from a CPE, before redirecting traffic into a legacy TCP session. Further Network-Assisted MPTCP deployment and operational considerations are discussed in [I-D.nam-mptcp-deployment-considerations].

```
  +------------+        _--------_    +----------------+
  |            |       (    LTE    )  |                |
  |   CPE      +=======+          +====+  Backbone      |
  |  (MCP)     |       (_         _)  |   Network      |
  |            |         (_______)    |+--------------+|
  |            |       IP Network #1  || Concentrator ||------> Internet
  |            |                      ||    (MCP)     ||
  |            |                      |+--------------+|
  |            |       IP Network #2  |                |
  |            |        _--------_    |                |
  |            |       (    DSL     ) |                |
  |            +=======+            +==+                |
  |            |       (_         _)  |                |
  +-----+------+        (_______)     +----------------+
        |
  ---- LAN ----
        |
    end-nodes
```

Figure 1: "Network-Assisted" MPTCP Design

This document focuses on the explicit mode that consists in configuring explicitly the reachability information of the MCP on a host. Concretely, the explicit mode has several advantages, e.g.,:

- o It does not impose any specific constraint on the location of the MCP. For example, the MCP can be located in any access network, located upstream in the core network, or located in a data canter facility.

- o Tasks required for activating the explicit mode are minimal. In particular, this mode does not require any specific routing and/or forwarding policies for handling outbound packets other than ensuring that an MCP is reachable from a CPE, and vice versa (which is straightforward IP routing policy operation).

o The engineering effort to change the location of an MCP for some reason (e.g., to better accommodate dimensioning constraints, to move the MCP to a data canter, to enable additional MCP instances closer to the customer premises, etc.) is minimal

o An operator can easily enforce strategies for differentiating the treatment of MPTCP connections that are directly initiated by an MPTCP-enabled host connected to an MCP if the explicit mode is enabled. Typically, an operator may decide to offload MPTCP connections originated by an MPTCP-enabled terminal from being forwarded through a specific MCP, or decide to relay them via a specific MCP. Such policies can be instructed to the MCP. Implementing such differentiating behavior if the implicit mode is in use may be complex to achieve.

o Multiple MCPs can be supported to service the same CPE, e.g., an MCP can be enabled for internal services (to optimize the delivery of some operator-specific services) while another MCP may be solicited for external services (e.g., access to the Internet). The explicit mode allows the deployment of such scenario owing to the provisioning of an MCP selection policy table that relies upon the destination IP prefixes to select the MCP to involve for an ongoing MPTCP connection, for instance.

o Because the MCP's reachability information is explicitly configured on the CPE, means to guarantee successful inbound connections can be enabled in the CPE to dynamically discover the external IP address that has been assigned for communicating with remote servers, instruct the MCP to maintain active bindings so that incoming packets can be successfully redirected towards the appropriate CPE, etc.

o Troubleshooting and root cause analysis may be facilitated in the explicit mode since faulty key nodes that may have caused a service degradation are known. Because of the loose adherence to the traffic forwarding and routing polices, troubleshooting a service degradation that is specific to multi-access serviced customers should first investigate the behavior of the involved MCP.

This document defines DHCPv4 [RFC2131] and DHCPv6 [RFC3315] options that can be used to configure hosts with MCP IP addresses.

This specification assumes an MCP is reachable through one or multiple IP addresses. As such, a list of IP addresses can be returned in the DHCP MPTCP option. Also, it assumes the various network attachments provided to an MPTCP-enabled CPE are managed by the same administrative entity.

2. Terminology

This document makes use of the following terms:

- o Multipath Conversion Point (MCP): a function that terminates a transport flow and relays all data received over it over another transport flow. This element is located upstream in the network. One or multiple MCPs can be deployed in the network side to assist MPTCP-enabled devices to establish MPTCP connections via available network attachments.

  On the uplink path, the MCP terminates the MPTCP connections [RFC6824] received from its customer-facing interfaces and transforms these connections into legacy TCP connections [RFC0793] towards upstream servers.

  On the downlink path, the MCP turns the legacy server's TCP connection into MPTCP connections towards its customer-facing interfaces.
- o DHCP refers to both DHCPv4 [RFC2131] and DHCPv6 [RFC3315].
- o DHCP client denotes a node that initiates requests to obtain configuration parameters from one or more DHCP servers.
- o DHCP server refers to a node that responds to requests from DHCP clients.

3. DHCPv6 MPTCP Option

3.1. Format

The DHCPv6 MPTCP option can be used to configure a list of IPv6 addresses of an MCP.

The format of this option is shown in Figure 2. As a reminder, this format follows the guidelines for creating new DHCPv6 options (Section 5.1 of [RFC7227]).

```
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |     OPTION_V6_MPTCP           |         Option-length         |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 |                         ipv6-address                          |
 |                                                               |
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 |                         ipv6-address                          |
 |                                                               |
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                              ...                              |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: DHCPv6 MPTCP option

The fields of the option shown in Figure 2 are as follows:

- Option-code: OPTION_V6_MPTCP (TBA, see Section 7.1)
- Option-length: Length of the 'MCP IP Address(es)' field in octets. MUST be a multiple of 16.
- MCP IPv6 Addresses: Includes one or more IPv6 addresses [RFC4291] of the MCP to be used by the MPTCP client.

  Note, IPv4-mapped IPv6 addresses (Section 2.5.5.2 of [RFC4291]) are allowed to be included in this option.

To return more than one MCPs to the requesting DHCPv6 client, the DHCPv6 server returns multiple instances of OPTION_V6_MPTCP. Some guidelines for DHCP servers are elaborated in Appendix A.

3.2. DHCPv6 Client Behavior

Clients MAY request option OPTION_V6_MPTCP, as defined in [RFC3315], Sections 17.1.1, 18.1.1, 18.1.3, 18.1.4, 18.1.5, and 22.7. As a convenience to the reader, we mention here that the client includes requested option codes in the Option Request Option.

The DHCPv6 client MUST be prepared to receive multiple instances of OPTION_V6_MPTCP; each instance is to be treated separately as it corresponds to a given MCP: there are as many MCPs as instances of the OPTION_V6_MPTCP option.

If an IPv4-mapped IPv6 address is received in OPTION_V6_MPTCP, it indicates that the MCP has the corresponding IPv4 address.

The DHCPv6 client MUST silently discard multicast and host loopback addresses [RFC6890] conveyed in OPTION_V6_MPTCP.

4. DHCPv4 MPTCP Option

4.1. Format

The DHCPv4 MPTCP option can be used to configure a list of IPv4 addresses of an MCP. The format of this option is illustrated in Figure 3.

```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Code         |     Length    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| List-Length   |   List of     |
+-+-+-+-+-+-+-+-+    MPTCP      |
/      MCP IPv4 Addresses       /
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+  ---
| List-Length   |   List of     |   |
+-+-+-+-+-+-+-+-+    MPTCP      |   |
/      MCP IPv4 Addresses       /   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+   |
.             ...               . Optional
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+   |
| List-Length   |   List of     |   |
+-+-+-+-+-+-+-+-+    MPTCP      |   |
/      MCP IPv4 Addresses       /   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+  ---
```

Figure 3: DHCPv4 MPTCP option

The fields of the option shown in Figure 3 are as follows:

- Code: OPTION_V4_MPTCP (TBA, see Section 7.2);
- Length: Length of all included data in octets. The minimum length is 5.
- List-Length: Length of the "List of MCP IPv4 Addresses" field in octets; MUST be a multiple of 4.
- List of MCP IPv4 Addresses: Contains one or more IPv4 addresses of the MCP to be used by the MPTCP client. The format of this field is shown in Figure 4.
- OPTION_V4_MPTCP can include multiple lists of MCP IPv4 addresses; each list is treated separately as it corresponds to a given MCP.

When several lists of MCP IPv4 addresses are to be included, "List-Length" and "MCP IPv4 Addresses" fields are repeated.

```
0     8     16    24    32    40    48
+-----+-----+-----+-----+-----+-----+--
|  a1 |  a2 |  a3 |  a4 |  a1 |  a2 | ...
+-----+-----+-----+-----+-----+-----+--
     IPv4 Address 1          IPv4 Address 2 ...
```

This format assumes that an IPv4 address is encoded as a1.a2.a3.a4.

Figure 4: Format of the List of MCP IPv4 Addresses

OPTION_V4_MPTCP is a concatenation-requiring option. As such, the mechanism specified in [RFC3396] MUST be used if OPTION_V4_MPTCP exceeds the maximum DHCPv4 option size of 255 octets.

Some guidelines for DHCP servers are elaborated in Appendix A.

4.2. DHCPv4 Client Behavior

To discover one or more MCPs, the DHCPv4 client MUST include OPTION_V4_MPTCP in a Parameter Request List Option [RFC2132].

The DHCPv4 client MUST be prepared to receive multiple lists of MCP IPv4 addresses in the same OPTION_V4_MPTCP; each list is to be treated as a separate MCP instance.

The DHCPv4 client MUST silently discard multicast and host loopback addresses [RFC6890] conveyed in OPTION_V4_MPTCP.

5. Security Considerations

The security considerations in [RFC2131] and [RFC3315] are to be considered.

MPTCP-related security considerations are discussed in [RFC6824].

Means to protect the MCP against Denial-of-Service (DoS) attacks must be enabled. Such means include the enforcement of ingress filtering policies at the boundaries of the network. In order to prevent exhausting the resources of the MCP by creating an aggressive number of simultaneous subflows for each MPTCP connection, the administrator should limit the number of allowed subflows per host for a given connection.

Attacks outside the domain can be prevented if ingress filtering is enforced. Nevertheless, attacks from within the network between a

host and an MCP instance are yet another actual threat. Means to ensure that illegitimate nodes cannot connect to a network should be implemented.

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

6. Privacy Considerations

Generic privacy-related considerations are discussed in [RFC7844].

The MCP may have access to privacy-related information (e.g., International Mobile Subscriber Identity (IMSI), link identifier, subscriber credentials, etc.). The MCP must not leak such sensitive information outside an administrative domain.

7. IANA Considerations

7.1. DHCPv6 Option

IANA is requested to assign the following new DHCPv6 Option Code in the registry maintained in http://www.iana.org/assignments/dhcpv6-parameters:

Option Name	Value
OPTION_V6_MPTCP	TBA

7.2. DHCPv4 Option

IANA is requested to assign the following new DHCPv4 Option Code in the registry maintained in http://www.iana.org/assignments/bootp-dhcp-parameters/:

Option Name	Value	Data length	Meaning
OPTION_V4_MPTCP	TBA	Variable; the minimum length is 5.	Includes one or multiple lists of MCP IP addresses; each list is treated as a separate MCP.

8. Acknowledgements

Many thanks to Olivier Bonaventure for the feedback on this document. Olivier suggested to define the option as a name but that design approach was debated several times within the dhc wg.

Thanks to Dan Seibel, Bernie Volz, Niall O'Reilly, Simon Hobson, and Ted Lemon for the feedback on the dhc wg mailing list.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <http://www.rfc-editor.org/info/rfc2131>.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, DOI 10.17487/RFC2132, March 1997, <http://www.rfc-editor.org/info/rfc2132>.

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <http://www.rfc-editor.org/info/rfc3315>.

[RFC3396] Lemon, T. and S. Cheshire, "Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4)", RFC 3396, DOI 10.17487/RFC3396, November 2002, <http://www.rfc-editor.org/info/rfc3396>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <http://www.rfc-editor.org/info/rfc4291>.

[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>.

[RFC6890] Cotton, M., Vegoda, L., Bonica, R., Ed., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, DOI 10.17487/RFC6890, April 2013, <http://www.rfc-editor.org/info/rfc6890>.

9.2. Informative References

[I-D.nam-mptcp-deployment-considerations] Boucadair, M., Jacquenet, C., Bonaventure, O., Henderickx, W., and R. Skog, "Network-Assisted MPTCP: Use Cases, Deployment Scenarios and Operational Considerations", draft-nam-mptcp-deployment-considerations-01 (work in progress), December 2016.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <http://www.rfc-editor.org/info/rfc793>.

[RFC7227] Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", BCP 187, RFC 7227, DOI 10.17487/RFC7227, May 2014, <http://www.rfc-editor.org/info/rfc7227>.

[RFC7844] Huitema, C., Mrugalski, T., and S. Krishnan, "Anonymity Profiles for DHCP Clients", RFC 7844, DOI 10.17487/RFC7844, May 2016, <http://www.rfc-editor.org/info/rfc7844>.

[RFC7969] Lemon, T. and T. Mrugalski, "Customizing DHCP Configuration on the Basis of Network Topology", RFC 7969, DOI 10.17487/RFC7969, October 2016, <http://www.rfc-editor.org/info/rfc7969>.

Appendix A. DHCP Server Configuration Guidelines

DHCP servers that support the DHCP MCP option can be configured with a list of IP addresses of the MCP(s). If multiple IP addresses are configured, the DHCP server MUST be explicitly configured whether all or some of these addresses refer to:

1. the same MCP: the DHCP server returns multiple addresses in the same instance of the DHCP MCP option.

2. distinct MCPs : the DHCP server returns multiple lists of MCP IP addresses to the requesting DHCP client (encoded as multiple OPTION_V6_MPTCP or in the same OPTION_V4_MPTCP); each list refers to a distinct MCP.

Precisely how DHCP servers are configured to separate lists of IP addresses according to which MCP they refer to is out of scope for this document. However, DHCP servers MUST NOT combine the IP addresses of multiple MCPs and return them to the DHCP client as if they were belonging to a single MCP, and DHCP servers MUST NOT separate the addresses of a single MCP and return them as if they were belonging to distinct MCPs. For example, if an administrator configures the DHCP server by providing a Fully Qualified Domain Name (FQDN) for an MCP, even if that FQDN resolves to multiple addresses, the DHCP server MUST deliver them within a single server address block.

DHCPv6 servers that implement this option and that can populate the option by resolving FQDNs will need a mechanism for indicating whether to query A records or only AAAA records. When a query returns A records, the IP addresses in those records are returned in the DHCPv6 response as IPv4-mapped IPv6 addresses.

Since this option requires support for IPv4-mapped IPv6 addresses, a DHCPv6 server implementation will not be complete if it does not query A records and represent any that are returned as IPv4-mapped IPv6 addresses in DHCPv6 responses. The mechanism whereby DHCPv6 implementations provide this functionality is beyond the scope of this document.

For guidelines on providing context-specific configuration information (e.g., returning a regional-based configuration), and information on how a DHCP server might be configured with FQDNs that get resolved on demand, see [RFC7969].

Authors' Addresses

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Christian Jacquenet
Orange
Rennes
France

Email: christian.jacquenet@orange.com

Tirumaleswar Reddy
Cisco Systems, Inc.
Cessna Business Park, Varthur Hobli
Sarjapur Marathalli Outer Ring Road
Bangalore, Karnataka 560103
India

Email: tireddy@cisco.com