

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 9, 2020

M. Boucadair
C. Jacquenet
Orange
T. Reddy
McAfee
October 7, 2019

DHCP Options for 0-RTT TCP Converters
draft-boucadair-tcpm-dhc-converter-03

Abstract

Because of the lack of important TCP extensions, e.g., Multipath TCP support at the server side, some service providers now consider a network-assisted model that relies upon the activation of a dedicated function called Transport Converters. For example, network-assisted Multipath TCP deployment models are designed to facilitate the adoption of Multipath TCP for the establishment of multi-path communications without making any assumption about the support of Multipath TCP by the remote servers. Transport Converters located in the network are responsible for establishing multi-path communications on behalf of endpoints, thereby taking advantage of Multipath TCP 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 Transport Converters is explicitly configured on connected hosts. This document specifies DHCP (IPv4 and IPv6) options to configure hosts with Converters parameters.

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."

Internet-Draft

DHCP for TCP Converters

October 2019

This Internet-Draft will expire on April 9, 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.	Terminology	3
3.	DHCPv6 Converter Option	4
3.1.	Format	4
3.2.	DHCPv6 Client Behavior	5
4.	DHCPv4 Converter Option	6
4.1.	Format	6
4.2.	DHCPv4 Client Behavior	7
5.	Security Considerations	7
6.	IANA Considerations	8
6.1.	DHCPv6 Option	8
6.2.	DHCPv4 Option	8
7.	Acknowledgements	8
8.	References	9
8.1.	Normative References	9
8.2.	Informative References	10
Appendix A.	DHCP Server Configuration Guidelines	10
	Authors' Addresses	11

[1.](#) Introduction

One of the promising deployment scenarios for Multipath TCP (MPTCP, [\[RFC6824\]](#)) is to enable a host or a Customer Premises Equipment (CPE) connected to multiple networks (e.g., DSL, LTE, WLAN) to optimize the

usage of such resources. A deployment scenario relies on MPTCP Conversion Points (called, Transport Converters [[I-D.ietf-tcpm-converters](#)]). A Converter terminates the extended TCP (e.g., MPTCP, TCPinc) sessions established from a host, 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](#)].

Figure 1 shows a deployment example of the Converters to assist establishing MPTCP connections.

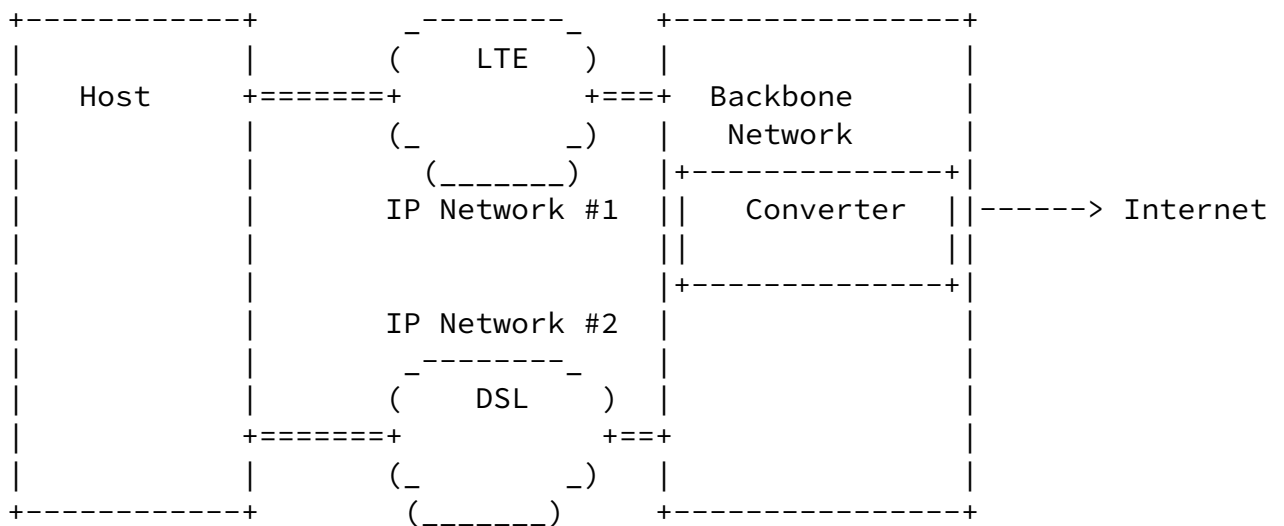


Figure 1: "Network-Assisted" MPTCP Design

[[I-D.ietf-tcpm-converters](#)] specifies the Converter as a function that is installed by a network operator to aid the deployment of TCP extensions and to provide the benefits of such extensions to clients. A Transport Converter supports one or more TCP extensions.

[[I-D.ietf-tcpm-converters](#)] assumes the explicit mode that consists in configuring explicitly the reachability information of the Converter(s) on a host.

This document defines DHCPv4 [[RFC2131](#)] and DHCPv6 [[RFC8415](#)] options that can be used to configure hosts with Converter IP addresses.

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

[2.](#) Terminology

This document makes use of the following terms:

- o Converter: 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 Converters can be deployed in the network side. The Converter achieves the following:

- * Listen for client sessions;
 - * Receive from a client the address of the final target server;
 - * Setup a session to the final server;
 - * Relay control messages and data between the client and the server;
 - * Perform access controls according to local policies.
- o DHCP refers to both DHCPv4 [[RFC2131](#)] and DHCPv6 [[RFC8415](#)].
 - 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.

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.

[3.](#) DHCPv6 Converter Option

[3.1.](#) Format

The DHCPv6 Converter option can be used to configure a list of IPv6 addresses of a Converter.

- o Option-code: OPTION_V6_CONVERT (TBA, see [Section 6.1](#))
- o Option-length: Length of the 'Converter IP Address(es)' field in octets. MUST be a multiple of 16.
- o Converter IPv6 Addresses: Includes one or more IPv6 addresses [[RFC4291](#)] of the Converter to be used by the TCP 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 Converter to the requesting DHCPv6 client, the DHCPv6 server returns multiple instances of OPTION_V6_CONVERT. Some guidelines for DHCP servers are elaborated in [Appendix A](#).

[3.2.](#) DHCPv6 Client Behavior

Clients MAY request option OPTION_V6_CONVERT, as defined in [[RFC8415](#)], Sections [18.2.1](#), [18.2.2](#), [18.2.4](#), [18.2.5](#), [18.2.6](#), and [21.7](#). 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_CONVERT; each instance is to be treated separately as it corresponds to a given Converter: there are as many Converters as instances of the OPTION_V6_CONVERT option.

If an IPv4-mapped IPv6 address is received in OPTION_V6_CONVERT, it indicates that the Converter has the corresponding IPv4 address.

The DHCPv6 client MUST silently discard multicast and host loopback addresses [[RFC6890](#)] conveyed in OPTION_V6_CONVERT.

[4.](#) DHCPv4 Converter Option

[4.1.](#) Format

The DHCPv4 Converter option can be used to configure a list of IPv4 addresses of a Converter. 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
```




Figure 3: DHCPv4 Converter option

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

- o Code: OPTION_V4_CONVERT (TBA, see [Section 6.2](#));
- o Length: Length of all included data in octets. The minimum length is 5.
- o List-Length: Length of the "List of Converter IPv4 Addresses" field in octets; MUST be a multiple of 4.
- o List of Converter IPv4 Addresses: Contains one or more IPv4 addresses of the Converter to be used by the TCP client. The format of this field is shown in Figure 4.

- o OPTION_V4_CONVERT can include multiple lists of Converter IPv4 addresses; each list is treated separately as it corresponds to a given Converter.

When several lists of Converter IPv4 addresses are to be included, "List-Length" and "Converter 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 Converter IPv4 Addresses

OPTION_V4_CONVERT is a concatenation-requiring option. As such, the mechanism specified in [\[RFC3396\]](#) MUST be used if OPTION_V4_CONVERT 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 Converters, the DHCPv4 client MUST include OPTION_V4_CONVERT in a Parameter Request List Option [\[RFC2132\]](#).

The DHCPv4 client MUST be prepared to receive multiple lists of Converter IPv4 addresses in the same OPTION_V4_CONVERT; each list is to be treated as a separate Converter instance.

The DHCPv4 client MUST silently discard multicast and host loopback addresses [\[RFC6890\]](#) conveyed in OPTION_V4_CONVERT.

[5](#). Security Considerations

The security considerations in [\[RFC2131\]](#) and [\[RFC8415\]](#) are to be considered.

Generic Convert security considerations are discussed in [\[I-D.ietf-tcpm-converters\]](#).

MPTCP-related security considerations are discussed in [\[RFC6824\]](#).

Means to protect the Converter 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 Converter 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 a Converter 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 Converter is inserted in the path. Indeed, inserting an illegitimate Converter 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 a Converter should be enabled.

[6.](#) IANA Considerations

[6.1.](#) 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_CONVERT	TBA

[6.2.](#) 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_CONVERT	TBA	Variable; 5.	Includes one or multiple lists of the minimum Converter IP addresses; each list is treated as a separate Converter.

[7.](#) 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 dhcp wg mailing list.

[8.](#) References

[8.1.](#) Normative References

- [I-D.ietf-tcpm-converters]
Bonaventure, O., Boucadair, M., Gundavelli, S., Seo, S., and B. Hesmans, "0-RTT TCP Convert Protocol", [draft-ietf-tcpm-converters-12](#) (work in progress), October 2019.
- [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>>.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", [RFC 2131](#), DOI 10.17487/RFC2131, March 1997, <<https://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, <<https://www.rfc-editor.org/info/rfc2132>>.
- [RFC3396] Lemon, T. and S. Cheshire, "Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4)", [RFC 3396](#), DOI 10.17487/RFC3396, November 2002, <<https://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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc6890>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174,

- [RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 8415](#), DOI 10.17487/RFC8415, November 2018, <<https://www.rfc-editor.org/info/rfc8415>>.

8.2. 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, <<https://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, <<https://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, <<https://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, <<https://www.rfc-editor.org/info/rfc7969>>.

Appendix A. DHCP Server Configuration Guidelines

DHCP servers that support the DHCP Converter option can be configured with a list of IP addresses of the Converter(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 Converter: the DHCP server returns multiple addresses in the same instance of the DHCP Converter option.
2. distinct Converters : the DHCP server returns multiple lists of Converter IP addresses to the requesting DHCP client (encoded as multiple OPTION_V6_CONVERT or in the same OPTION_V4_CONVERT); each list refers to a distinct Converter.

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

Mohamed Boucadair
Orange
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Christian Jacquenet
Orange
Rennes
France

Email: christian.jacquenet@orange.com

Boucadair, et al.

Expires April 9, 2020

[Page 11]

Internet-Draft

DHCP for TCP Converters

October 2019

Tirumaleswar Reddy
McAfee, Inc.
Embassy Golf Link Business Park
Bangalore, Karnataka 560071
India

Email: kondtir@gmail.com

