PCP Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 23, 2013

M. Boucadair
France Telecom
R. Penno
D. Wing
Cisco

February 19, 2013

DHCP Options for the Port Control Protocol (PCP)
draft-ietf-pcp-dhcp-06

Abstract

This document specifies DHCP (IPv4 and IPv6) options to configure hosts with Port Control Protocol (PCP) Server names. The use of DHCPv4 or DHCPv6 depends on the PCP deployment scenario.

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 August 23, 2013.

Copyright Notice

Table of Contents

1. Introduction 3
2. Terminology 3
3. DHCPv6 PCP Server Option 3
3.1. Format 3
3.2. Client Behavior 4
4. DHCPv4 PCP Option 5
4.1. Format 5
4.2. Client Behavior 6
5. Use of PCP Server Names 6
6. Dual-Stack Hosts 7
7. Security Considerations 7
8. IANA Considerations 7
8.1. DHCPv6 Option 7
8.2. DHCPv4 Option 7
9. Acknowledgements 7
10. References 8
10.1. Normative References 8
10.2. Informative References 8
Appendix A. Rationale 9
A.1. Dependency on Name Resolution 10
Authors' Addresses 11

1. Introduction

This document defines DHCPv4 [RFC2131] and DHCPv6 [RFC3315] options which can be used to provision PCP Server [I-D.ietf-pcp-base] names. Motivations for expressing the PCP option as a textual string rather than a 32 or 128-bit binary address are discussed in Appendix A.

In order to make use of these options, this document assumes appropriate name resolution means (e.g., Section 6.1.1 of [RFC1123]) are available on the host client.

The use of DHCPv4 or DHCPv6 depends on the PCP deployment scenario.

2. Terminology

This document makes use of the following terms:

- PCP Server denotes a functional element which receives and processes PCP requests from a PCP Client. A PCP Server can be co-located with or be separated from the function (e.g., NAT, Firewall) it controls. Refer to [I-D.ietf-pcp-base].
- PCP Client denotes a PCP software instance responsible for issuing PCP requests to a PCP Server. Refer to [I-D.ietf-pcp-base].
- DHCP refers to both DHCPv4 [RFC2131] and DHCPv6 [RFC3315].
- DHCP client (or client) denotes a node that initiates requests to obtain configuration parameters from one or more DHCP servers.
- DHCP server (or server) refers to a node that responds to requests from DHCP clients.
- Name is a UTF-8 [RFC3629] string that can be passed to getaddrinfo (Section 6.1 of [RFC3493]), such as a DNS name, address literals, etc. The name MUST NOT contain spaces or nulls. A name may be a fully qualified domain name (e.g., "myservice.example.com."), IPv4 address in dotted-decimal form (e.g., 192.0.2.33) or textual representation of an IPv6 address (e.g., 2001:db8::1) [RFC4291][RFC5952].

3. DHCPv6 PCP Server Option

This DHCPv6 option conveys a name to be used to retrieve the IP addresses of PCP Server(s). Appropriate name resolution queries should be issued to resolve the conveyed name.

3.1. Format

The format of the DHCPv6 PCP Server option is shown in Figure 1.

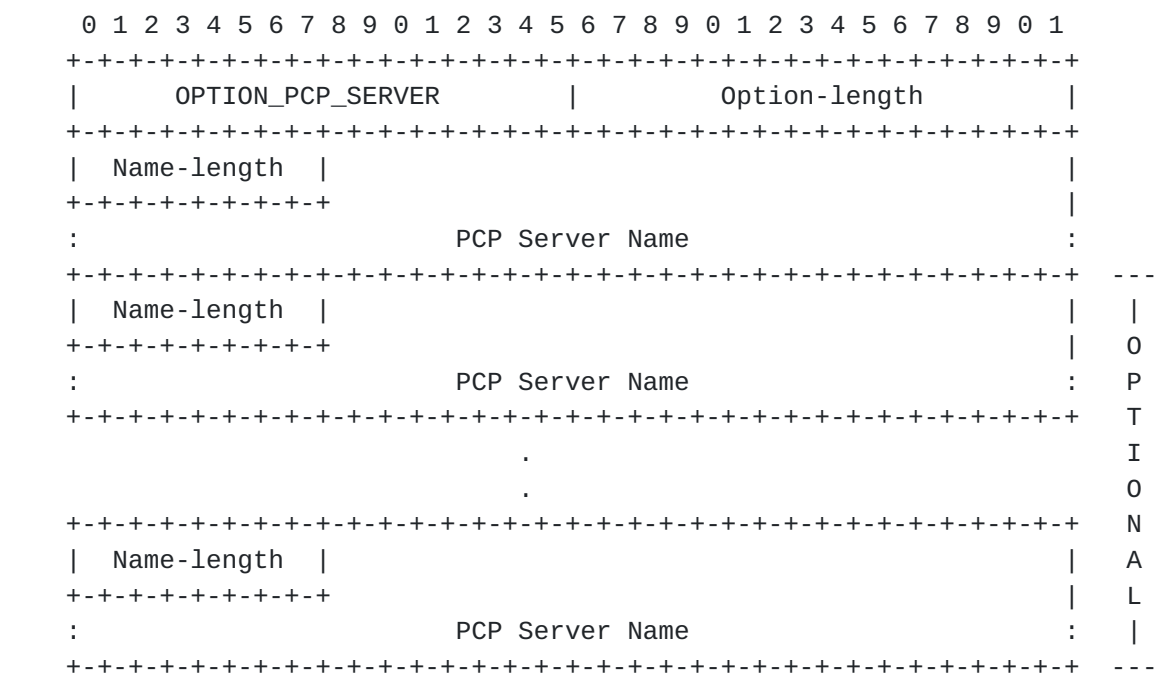

Figure 1: PCP Server Name DHCPv6 Option

The fields of the option shown in Figure 1 are as follows:

- Option-code: OPTION_PCP_SERVER (TBA, see Section 8.1)
- Option-length: includes total length of all following option data in octets.
- Name-length (one-octet field): Includes the length of the PCP Server Name, in octets.
- PCP Server Name (variable): The name of the PCP Server to be used by the PCP Client. The name is encoded as a UTF-8 [RFC3629] string.

The OPTION_PCP_SERVER option can include multiple PCP Server names; each name is treated as a separate PCP Server. When several names are to be included, "Name-length" and "PCP Server Name" fields are repeated.

3.2. Client Behavior

To discover a PCP Server [I-D.ietf-pcp-base], the DHCPv6 client MUST include an Option Request Option (ORO) requesting the DHCPv6 PCP Server Name option as described in Section 22.7 of [RFC3315] (i.e., include OPTION_PCP_SERVER on its OPTION_ORO).

If the DHCPv6 client receives an OPTION_PCP_SERVER option from the DHCPv6 server, it extracts the name(s) conveyed in the OPTION_PCP_SERVER option. A name is considered as valid if it is a

legal UTF-8 string which does not contain any spaces or nulls. Below are listed some additional validation rules:

- The trailing dot is optional when a domain name is conveyed in the option.
- IPv6 addresses MUST NOT be enclosed in brackets.
- A domain name is structured as one or more labels concatenated with dots. A label MUST have no more than 63 characters.

The DHCPv6 client MUST silently ignore invalid names.

Once each name conveyed in the OPTION_PCP_SERVER option is validated, the DHCPv6 client MUST follow the procedure specified in Section 5.

4. DHCPv4 PCP Option

4.1. Format

The PCP Server Name DHCPv4 option can be used to configure a name to be used by the PCP Client to contact a PCP Server. The format of this option is illustrated in Figure 2.

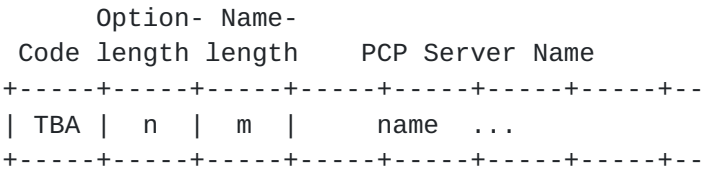

Figure 2: PCP Server Name DHCPv4 Option

The description of the fields is as follows:

- Code: OPTION_PCP_SERVER (TBA, see Section 8.2);
- Option-length: includes total length of all following option data in octets. The maximum length is 255 octets.
- Name-length (one-octet field): Includes the length of the PCP Server Name, in octets.
- PCP Server Name (variable): The name of the PCP Server to be used by the PCP Client when issuing PCP messages. The name is encoded as a UTF-8 [RFC3629] string.

The OPTION_PCP_SERVER option can include multiple PCP Server names; each name is treated as a separate PCP Server. When several names are to be included, "Name-length" and "PCP Server Name" fields are repeated.

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

4.2. Client Behavior

DHCPv4 client expresses the intent to get OPTION_PCP_SERVER by specifying it in Parameter Request List Option [RFC2132].

If the DHCPv4 client receives an OPTION_PCP_SERVER option from the DHCPv4 server, it extracts the name(s) conveyed in the option. A name is considered as valid if it is a legal UTF-8 string which does not contain any spaces or nulls. Below are listed some additional validation rules:

- The trailing dot is optional when a domain name is conveyed in the option.

- A domain name is structured as one or more labels concatenated with dots. A label MUST have no more than 63 characters.

The DHCPv4 client MUST silently discard non valid names.

Once each name conveyed in the OPTION_PCP_SERVER option is validated, the DHCPv4 client MUST follow the procedure specified in Section 5.

5. Use of PCP Server Names

Each configured PCP Server Name is passed to the name resolution library (e.g., Section 6.1.1 of [RFC1123] or [RFC6055]) to retrieve the corresponding IP address(es) (IPv4 or IPv6). It is out of scope of this document to specify how the PCP Client selects the PCP Server(s) to contact.

Multiple PCP Server Names may be configured to a PCP Client in some deployment contexts such as multi-homing. It is out of scope of this document to enumerate all deployment scenarios which require multiple Names to be configured.

A host may have multiple network interfaces (e.g, 3G, WiFi, etc.); each configured differently. Each PCP Server learned MUST be associated with the interface via which it was learned.

6. Dual-Stack Hosts

In some deployment contexts, the PCP Server may be reachable with an IPv4 address but DHCPv6 is used to provision the PCP Client. In such scenarios, a plain IPv4 address or an IPv4-mapped IPv6 address can be configured to reach the PCP Server.

A Dual-Stack host may receive OPTION_PCP_SERVER via both DHCPv4 and DHCPv6. The content of these OPTION_PCP_SERVER options may refer to the same or distinct PCP Servers. This is deployment-specific and as such it is out of scope of this document.

7. Security Considerations

The security considerations in [RFC2131], [RFC3315] and [I-D.ietf-pcp-base] are to be considered.

8. IANA Considerations

8.1. DHCPv6 Option

Authors of this document request the following DHCPv6 option code:

Option Name	Value
OPTION_PCP_SERVER	TBA

8.2. DHCPv4 Option

Authors of this document request the following DHCPv4 option code:

Option Name	Value
OPTION_PCP_SERVER	TBA

9. Acknowledgements

Many thanks to B. Volz, C. Jacquenet, R. Maglione, D. Thaler, T. Mrugalski, T. Lemon, S. Cheshire and M. Wasserman for their review and comments.

10. References

10.1. Normative References

[I-D.ietf-pcp-base]
Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", draft-ietf-pcp-base-29 (work in progress), November 2012.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC3396] Lemon, T. and S. Cheshire, "Encoding Long Options in the Dynamic Host Configuration Protocol (DHCPv4)", RFC 3396, November 2002.

[RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

[RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010.

10.2. Informative References

[I-D.ietf-behave-lsn-requirements]
Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common requirements for Carrier Grade NATs (CGNs)", draft-ietf-behave-lsn-requirements-10 (work in progress), December 2012.

[I-D.ietf-dhc-option-guidelines]
Hankins, D., Mrugalski, T., Siodelski, M., Jiang, S., and S. Krishnan, "Guidelines for Creating New DHCPv6 Options", draft-ietf-dhc-option-guidelines-09 (work in progress), December 2012.

[RFC1123] Braden, R., "Requirements for Internet Hosts - Application

and Support", STD 3, RFC 1123, October 1989.

[RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, July 1997.

[RFC3493] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W. Stevens, "Basic Socket Interface Extensions for IPv6", RFC 3493, February 2003.

[RFC6055] Thaler, D., Klensin, J., and S. Cheshire, "IAB Thoughts on Encodings for Internationalized Domain Names", RFC 6055, February 2011.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, August 2011.

[RFC6334] Hankins, D. and T. Mrugalski, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite", RFC 6334, August 2011.

Appendix A. Rationale

The pcp WG consensus is to define a DHCP option which contains a string that can be passed to APIs (e.g., getaddrinfo()). In particular, the option should be designed to include a name (e.g., domain name [RFC2181]) or IP address literal string. In such design, DHCP clients are expected to pass the conveyed string to any supported name resolution library (DNS is a name resolution service among others). The underlying name resolution library is responsible for validating the name.

Distinct IP-Address and Name DHCP options have been considered in early stages of this specification. This flexibility aims to let service providers make their own engineering choices and use the most convenient option according to their deployment context. Nevertheless, the DHC WG's position is this flexibility has some drawbacks such as inducing errors (See Section 7 of [I-D.ietf-dhc-option-guidelines]). Therefore, only the Name option is maintained within this document.

This choice is motivated by operational considerations: In particular, some Service Providers are considering two levels of

redirection:

(1) The first level is national-wise and undertaken by DHCP: a regional-specific Name will be returned;
(2) The second level is done during the resolution of the regional-specific Name to redirect the customer to a regional PCP server among a pool deployed regionally.

Distinct operational teams are responsible for each of the above mentioned levels. A clear separation between the functional perimeter of each team is a sensitive task for the maintenance of the offered services. Regional teams will require to introduce new resources (e.g., new PCP-controlled devices such as Carrier Grade NATs (CGNs, [I-D.ietf-behave-lsn-requirements])) to meet an increase in customer base. Operations related to the introduction of these new devices (e.g., addressing, redirection, etc.) are implemented locally. Having this regional separation provides flexibility to manage portions of network operated by dedicated teams. This two-level redirection can not be met by the IP Address option.

In addition to the operational considerations:

- The use of the Name for NAT64 [RFC6146] might be suitable for load-balancing purposes;
- For the DS-Lite case [RFC6333], if the encapsulation mode is used to send PCP messages, an IP address may be used since the AFTR selection is already done via the AFTR_NAME DHCPv6 option [RFC6334]. Of course, this assumes that the PCP Server is co-located with the AFTR function. If these functions are not co-located, conveying the Name would be more convenient.

A.1. Dependency on Name Resolution

The approach adopted in this document allows for an IP address or a Name to be returned in the specified DHCP option. In particular, a server can resolve first the name and return in the option the resolved IP address(es). For deployments where this is not possible, the server can return a name which will be resolved by the host embedding the client. This document does not have any requirement on the underlying name resolution library (in particular, DNS is not assumed as the only available name resolution service).

Returning a Name requires the host to embed a name resolution service. Some may present this as an argument against defining a Name option. Nevertheless, this argument may be objected as implementing a name resolution library (e.g., embed a DNS resolver) is cheap and devices which don't embed DNS resolver are uncommon.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes, 35000
France

Email: mohamed.boucadair@orange.com

Reinaldo Penno
Cisco
USA

Email: repenno@cisco.com

Dan Wing
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134
USA

Email: dwing@cisco.com