

PCP Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 3, 2011

M. Boucadair
France Telecom
R. Penno
Juniper Networks
D. Wing
Cisco
March 2, 2011

DHCP and DHCPv6 Options for Port Control Protocol (PCP)
draft-bpw-pcp-dhcp-03

Abstract

This document specifies DHCP (IPv4 and IPv6) options to provision Port Control Protocol (PCP) Servers. The use of DHCP 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 September 3, 2011.

Copyright Notice

Copyright (c) 2011 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. Rationale	3
4. Consistent NAT and PCP Configuration	4
5. DHCPv6 PCP Server Option	5
5.1. Format	5
5.2. Client Behaviour	5
5.3. Server Behaviour	6
6. DHCP PCP Option	6
7. Security Considerations	8
8. IANA Considerations	8
9. Acknowledgements	9
10. References	9
10.1. Normative References	9
10.2. Informative References	9
Authors' Addresses	10

1. Introduction

This document defines DHCP [RFC2131] and DHCPv6 [RFC3315] options which can be used to provision a PCP Server [I-D.ietf-pcp-base] reachability information; more precisely it defines DHCP options to convey a FQDN (as per Section 3.1 of [RFC1035]) of a PCP Server. In order to make use of these options, this document assumes that a DNS server is configured on the host client by DHCP or other means.

The use of DHCP or DHCPv6 depends on the PCP deployment scenarios.

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. As described in Section 3.1 of [I-D.wing-behave-dns64-config], dual-stack hosts can issue IPv4 datagrams successfully to that IP address.

2. Terminology

This document makes use of the following terms:

- o PCP Server: A functional element which receives and processes PCP requests from a PCP Client. A PCP Server can be co-located or be separated from the function (e.g., NAT, Firewall) it controls. Refer to [I-D.ietf-pcp-base].
- o PCP Client: a PCP software instance responsible for issuing PCP requests to a PCP Server. Refer to [I-D.ietf-pcp-base].
- o DHCP client (or client) denotes a node that initiates requests to obtain configuration parameters from one or more DHCP servers [RFC3315].
- o DHCP server (or server) refers to a node that responds to requests from DHCP clients [RFC3315].

3. Rationale

[[Note: May be removed in future version of the I-D.]]

Both IP address and FQDN DHCP options have been defined in previous versions of this document. This flexibility aims to let service providers to make their own engineering choices and use the convenient option according to their deployment context. Nevertheless, DHC WG's position is this flexibility have some

drawbacks such as inducing errors. Therefore, only the FQDN option is maintained in this updated version.

This choice of defining the FQDN option rather than the IP address is motivated by operational considerations: In particular, some service providers are considering two levels of redirection: (1) The first level is national-wise is undertaken by DHCP: a regional-specific FQDN will be returned; (2) The second level is done during the resolution of the regional-specific FQDN to redirect the customer to a regional PCP Servers 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 CGNs) to meet an increase of customer base. The introduction of these new devices (addressing, redirection, etc.) is 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:

- o The use of the FQDN for NAT64 might be suitable for load-balancing purposes;
- o For the DS-Lite case, if the encapsulation mode is used to send PCP messages, an IP address may be used (the ipv4-mapped IPv6 address) since the AFTR selection is already done via the AFTR_NAME DHCPv6 option. Of course, this assumes that the PCP Server is co-located with the AFTR function. If these functions are not co-located, conveying the FQDN would be more convenient.

If the PCP Server is located in a LAN, a simple FQDN such as "pcp-server.local" can be used.

4. Consistent NAT and PCP Configuration

The PCP Server discovered through DHCP MUST be able to install mappings on the appropriate upstream PCP-controlled device that will be crossed by packets issued by the host or any terminal belonging to the same realm (e.g., DHCP client is embedded in a CP router). If DHCP is used to provision the FQDN of a PCP Server, an operator SHOULD configure appropriately DHCP servers to meet this requirement. In case this prerequisite is not met, customers would experience service troubles and their service(s) won't be delivered appropriately.

Note that this constraint is implicitly met in scenarios where only one single PCP-controlled device is deployed in the network.

5. DHCPv6 PCP Server Option

This DHCPv6 option conveys a domain name to be used to retrieve the IP address of a PCP Server. Appropriate DNS queries should be issued to resolve the conveyed FQDN. In the context of a DS-Lite architecture [I-D.ietf-softwire-dual-stack-lite], the retrieved address may be an IPv4 address or an IPv4-mapped IPv6 address [RFC4291] of a PCP Server, and in the case of NAT64 [I-D.ietf-behave-v6v4-xlate-stateful] an IPv6 address can be retrieved.

5.1. Format

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




Figure 1: PCP Server FQDN DHCPv6 Option

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

- o Option-code: OPTION_PCP_SERVER (TBA, see Section 8)
- o Option-length: Length of the 'PCP Server Domain Name' field in octets.
- o PCP Server Domain Name: The domain name of the PCP Server to be used by the PCP Client. The domain name is encoded as specified in Section 8 of [RFC3315].

5.2. Client Behaviour

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 FQDN option (i.e., include OPTION_PCP_SERVER on its OPTION_ORO) as described in Section 22.7 of [RFC3315]. A client MAY also include the OPTION_DNS_SERVERS option on its OPTION_ORO to

retrieve a DNS servers list.

If the DHCPv6 client receives more than one OPTION_PCP_SERVER option from the DHCPv6 server, it MUST discard all instances of that option.

Upon receipt of an OPTION_PCP_SERVER option, the DHCPv6 client MUST verify that the option length (including the Option-length octet) does not exceed 256 octets [RFC1035]). The DHCPv6 client MUST verify the FQDN is a properly encoded as detailed in Section 8 of [RFC3315].

Once the FQDN conveyed in a OPTION_PCP_SERVER option is validated, the DHCPv6 client issues appropriate DNS Query messages using the provided FQDN to resolve a AAAA Resource Record or A RR Resource Record. AAAA and A Queries may be issued in parallel or sequentially (e.g., using [RFC3484] selection process).

If the DNS response contains more than one IPv6/IPv4 address, the PCP Client MUST use the procedure defined in [I-D.ietf-pcp-base] for the address selection. The PCP Client MUST NOT send PCP requests to more than one IP address at the same time.

It is RECOMMENDED to associate a TTL with any address resulting from resolving the FQDN conveyed in a OPTION_PCP_SERVER DHCPv6 option when stored in a local cache. Considerations on how to flush out a local cache are out of the scope of this document.

5.3. Server Behaviour

A DHCPv6 server MUST NOT reply with a value for the OPTION_PCP_SERVER if the DHCPv6 client has not explicitly included OPTION_PCP_SERVER in its OPTION_ORO.

If OPTION_PCP_SERVER option is requested by the DHCPv6 client, DHCPv6 server MUST NOT send more than one OPTION_PCP_SERVER option in the response. The DHCPv6 server MUST include only one FQDN in a OPTION_PCP_SERVER option. The DHCPv6 server MUST NOT include an FQDN having a length exceeding 256 octets.

6. DHCP PCP Option

The PCP Server DHCP option can be used to configure a FQDN to be used by the PCP Client to contact a PCP Server. The generic format of this option is illustrated in Figure 2.

Because of the depletion of DHCP option codes and in order to anticipate future PCP-related DHCP options, the proposed option uses a sub-option field.

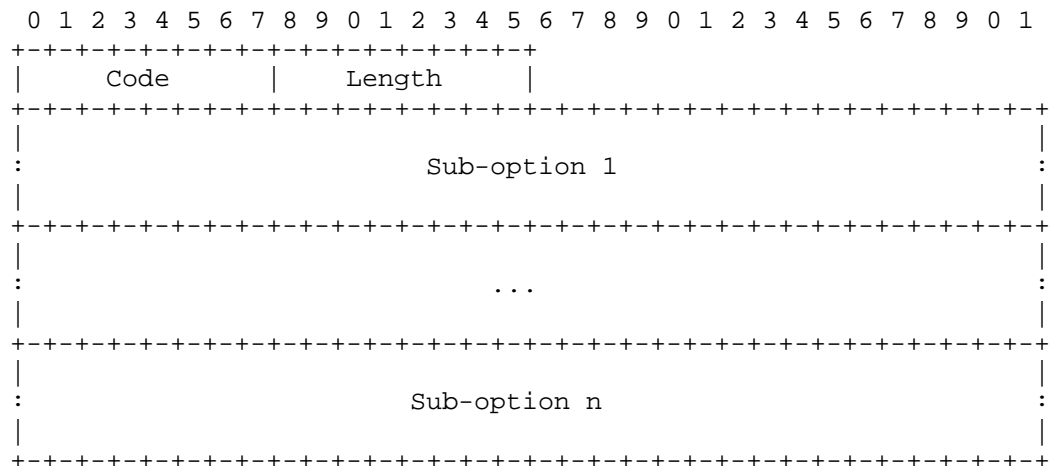


Figure 2: DHCP PCP Option

The description of the fields is as follows:

- o Code: OPTION_PCP_SERVER (TBA, see Section 8);
- o Length: Includes the length of included sub-options in octets; The maximum length is 255 octets.
- o One or several sub-options can be included in a PCP DHCP option. The format of each sub-option follows the structure shown in Figure 3.




Figure 3: PCP Server sub-option

Only one sub-option is defined in this document:

- 1: PCP Server Domain Name Sub-option (OPTION_PCP_SERVER_D (Figure 4)). This sub-option includes a FQDN of the PCP Server to be used by the PCP Client when issuing PCP messages.

Sub-option								
Code	Len	FQDN of PCP Server						
1	n	s1	s2	s3	s4	s5	...	

Figure 4: PCP Server FQDN DHCP Sub-option

The fields of the PCP Server Domain Name sub-option shown in Figure 4 are:

- o Sub-option Code: 1.
- o Len: Length of the "PCP Server Domain Name" field in octets.
- o PCP Server Domain Name: The domain name of the PCP Server to be used by the PCP Client. The encoding of the domain name is described in Section 3.1 of [RFC1035].

When the PCP Server Domain Name Sub-option is used, the client issues a DNS A record query to retrieve the IPv4 address(es) of the PCP server(s).

[[Note: aside effect of having the sub-option format is the risk to have a large option exceeding the maximum permissible within a single option (254 octets+the length octets). A solution would be to recommend [RFC3396]?]]

7. Security Considerations

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

8. IANA Considerations

This document requests the following codes:

DHCPv6 option code:

- o OPTION_PCP_SERVER

DHCP option code:

- o OPTION_PCP_SERVER

9. Acknowledgements

Many thanks to B. Volz, C. Jacquenet and R. Maglione 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 F. Dupont, "Port Control Protocol (PCP)",
draft-ietf-pcp-base-06 (work in progress), February 2011.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [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.
- [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.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

10.2. Informative References

- [I-D.ietf-behave-v6v4-xlate-stateful]
Bagnulo, M., Matthews, P., and I. Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers",
draft-ietf-behave-v6v4-xlate-stateful-12 (work in progress), July 2010.
- [I-D.ietf-softwire-dual-stack-lite]

Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", draft-ietf-softwire-dual-stack-lite-06 (work in progress), August 2010.

[I-D.wing-behave-dns64-config]

Wing, D., "IPv6-only and Dual Stack Hosts on the Same Network with DNS64", draft-wing-behave-dns64-config-03 (work in progress), February 2011.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes, 35000
France

Email: mohamed.boucadair@orange-ftgroup.com

Reinaldo Penno
Juniper Networks
1194 N Mathilda Avenue
Sunnyvale, California 94089
USA

Email: rpenno@juniper.net

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

Email: dwing@cisco.com

