

Internet Engineering Task Force (IETF)
Request for Comments: 7488
Updates: 6887
Category: Standards Track
ISSN: 2070-1721

M. Boucadair
France Telecom
R. Penno
D. Wing
P. Patil
T. Reddy
Cisco
March 2015

Port Control Protocol (PCP) Server Selection

Abstract

This document specifies the behavior to be followed by a Port Control Protocol (PCP) client to contact its PCP server(s) when one or several PCP server IP addresses are configured.

This document updates RFC 6887.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7488>.

Copyright Notice

Copyright (c) 2015 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	2
2. Terminology and Conventions	3
2.1. Requirements Language	3
2.2. Terminology	3
3. IP Address Selection: PCP Server with Multiple IP Addresses .	3
4. IP Address Selection: Multiple PCP Servers	4
5. Example: Multiple PCP Servers on a Single Interface	5
6. Security Considerations	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8
Appendix A. Multihoming	9
A.1. IPv6 Multihoming	9
A.2. IPv4 Multihoming	10
Acknowledgements	11
Authors' Addresses	12

1. Introduction

A host may have multiple network interfaces (e.g., 3G, IEEE 802.11, etc.), each configured with different PCP servers. Each PCP server learned must be associated with the interface on which it was learned. Generic multi-interface considerations are documented in Section 8.4 of [RFC6887]. Multiple PCP server IP addresses may be configured on a PCP client in some deployment contexts such as multihoming (see Appendix A). A PCP server may also have multiple IP addresses associated with it. It is out of the scope of this document to enumerate all deployment scenarios that require multiple PCP server IP addresses to be configured.

If a PCP client discovers multiple PCP server IP addresses, it needs to determine which actions it needs to undertake (e.g., whether PCP entries are to be installed in all or a subset of discovered IP addresses, whether some PCP entries are to be removed, etc.). This document makes the following assumptions:

- o There is no requirement that multiple PCP servers configured on the same interface have the same capabilities.
- o PCP requests to different PCP servers are independent, the result of a PCP request to one PCP server does not influence another.
- o The configuration mechanism must distinguish IP addresses that belong to the same PCP server.

This document specifies the behavior to be followed by a PCP client [RFC6887] to contact its PCP server(s) [RFC6887] when it is configured with one or several PCP server IP addresses (e.g., using DHCP [RFC7291]). This document does not make any assumption on the type of these IP addresses (i.e., unicast/anycast).

2. Terminology and Conventions

2.1. 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].

2.2. Terminology

- o PCP client: denotes a PCP software instance responsible for issuing PCP requests to a PCP server. Refer to [RFC6887].
- o PCP server: denotes a software instance that receives and processes PCP requests from a PCP client. A PCP server can be co-located with or be separated from the function it controls (e.g., Network Address Translation (NAT) or firewall). Refer to [RFC6887].

3. IP Address Selection: PCP Server with Multiple IP Addresses

This section describes the behavior a PCP client follows to contact its PCP server when the PCP client has multiple IP addresses for a single PCP server.

1. A PCP client should construct a set of candidate source addresses (see Section 4 of [RFC6724]) based on application input and PCP [RFC6887] constraints. For example, when sending a PEER or a MAP with a FILTER request for an existing TCP connection, the only candidate source address is the source address used for the existing TCP connection. But when sending a MAP request for a service that will accept incoming connections, the candidate source addresses may be all of the node's IP addresses or some subset of IP addresses on which the service is configured to listen.
2. The PCP client then sorts the PCP server IP addresses as per Section 6 of [RFC6724] using the candidate source addresses selected in the previous step as input to the destination address selection algorithm.

3. The PCP client initializes its Maximum Retransmission Count (MRC) to 4.
4. The PCP client sends its PCP messages following the retransmission procedure specified in Section 8.1.1 of [RFC6887]. If no response is received after MRC attempts, the PCP client retries the procedure with the next IP address in the sorted list.

The PCP client may receive a response from an IP address after exhausting MRC attempts for that particular IP address. The PCP client SHOULD ignore such a response because receiving a delayed response after exhausting four retransmissions sent with exponentially increasing intervals is an indication that problems are to be encountered in the corresponding forwarding path and/or when processing subsequent requests by that PCP server instance.

If, when sending PCP requests, the PCP client receives a hard ICMP error [RFC1122], it MUST immediately try the next IP address from the list of PCP server IP addresses.

5. If the PCP client has exhausted all IP addresses configured for a given PCP server, the procedure SHOULD be repeated every 15 minutes until the PCP request is successfully answered.
6. Once the PCP client has successfully received a response from a PCP server's IP address, all subsequent PCP requests to that PCP server are sent on the same IP address until that IP address becomes unresponsive. In case the IP address becomes unresponsive, the PCP client clears the cache of sorted destination addresses and follows the steps described above to contact the PCP server again.

For efficiency, the PCP client SHOULD use the same Mapping Nonce for requests sent to all IP addresses belonging to the same PCP server. As a reminder, nonce validation checks are performed when operating in the Simple Threat Model (see Section 18.1 of [RFC6887]) to defend against some off-path attacks.

4. IP Address Selection: Multiple PCP Servers

This section describes the behavior a PCP client follows to contact multiple PCP servers, with each PCP server reachable on a list of IP addresses. There is no requirement that these multiple PCP servers have the same capabilities.

Note that how PCP clients are configured to separate lists of IP addresses of each PCP server is implementation specific and deployment specific. For example, a PCP client can be configured using DHCP with multiple lists of PCP server IP addresses; each list is referring to a distinct PCP server [RFC7291].

If several PCP servers are configured, each with multiple IP addresses, the PCP client contacts all PCP servers using the procedure described in Section 3.


As specified in Sections 11.2 and 12.2 of [RFC6887], the PCP client must use a different Mapping Nonce for each PCP server with which it communicates.

If the PCP client is configured, using some means, with the capabilities of each PCP server, a PCP client may choose to contact all PCP servers simultaneously or iterate through them with a delay.

This procedure may result in a PCP client instantiating multiple mappings maintained by distinct PCP servers. The decision to use all these mappings or delete some of them depends on the purpose of the PCP request. For example, if the PCP servers are configuring firewall (not NAT) functionality, then the client would, by default (i.e., unless it knows that they all replicate state among them), need to use all the PCP servers.

5. Example: Multiple PCP Servers on a Single Interface

Figure 1 depicts an example that is used to illustrate the server selection procedure specified in Sections 3 and 4. In this example, PCP servers (A and B) are co-located with edge routers (rtr1 and rtr2) with each PCP server controlling its own device.



Edge Routers (rtr1, rtr2)

Figure 1: Single Uplink, Multiple PCP Servers

The example describes behavior when a single IP address for one PCP server is not responsive. The PCP client is configured with two PCP servers for the same interface, PCP-Server-A and PCP-Server-B, each of which have two IP addresses: an IPv4 address and an IPv6 address. The PCP client wants an IPv4 mapping, so it orders the addresses as follows:

- o PCP-Server-A:
 - * 192.0.2.1
 - * 2001:db8:1111::1
- o PCP-Server-B:
 - * 198.51.100.1
 - * 2001:db8:2222::1

Suppose that:

- o The path to reach 192.0.2.1 is broken
- o The path to reach 2001:db8:1111::1 is working
- o The path to reach 198.51.100.1 is working
- o The path to reach 2001:db8:2222::1 is working

It sends two PCP requests at the same time, the first to 192.0.2.1 (corresponding to PCP-Server-A) and the second to 198.51.100.1 (corresponding to PCP-Server-B). The path to 198.51.100.1 is working, so a PCP response is received. Because the path to 192.0.2.1 is broken, no PCP response is received. The PCP client retries four times to elicit a response from 192.0.2.1 and finally gives up on that address and sends a PCP message to 2001::db8:1111:1. That path is working, and a response is received. Thereafter, the PCP client should continue using that responsive IP address for PCP-Server-A (2001:db8:1111::1). In this particular case, it will have to use the THIRD_PARTY option for IPv4 mappings.

6. Security Considerations

PCP-related security considerations are discussed in [RFC6887].

This document does not specify how PCP server addresses are provisioned on the PCP client. It is the responsibility of PCP server provisioning document(s) to elaborate on security considerations to discover legitimate PCP servers.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012, <<http://www.rfc-editor.org/info/rfc6724>>.
- [RFC6887] Wing, D., Ed., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013, <<http://www.rfc-editor.org/info/rfc6887>>.

7.2. Informative References

- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989, <<http://www.rfc-editor.org/info/rfc1122>>.
- [RFC4116] Abley, J., Lindqvist, K., Davies, E., Black, B., and V. Gill, "IPv4 Multihoming Practices and Limitations", RFC 4116, July 2005, <<http://www.rfc-editor.org/info/rfc4116>>.
- [RFC7291] Boucadair, M., Penno, R., and D. Wing, "DHCP Options for the Port Control Protocol (PCP)", RFC 7291, July 2014, <<http://www.rfc-editor.org/info/rfc7291>>.

Appendix A. Multihoming

The main problem of a PCP multihoming situation can be succinctly described as "one PCP client, multiple PCP servers." As described in Section 3, if a PCP client discovers multiple PCP servers, it should send requests to all of them with assumptions described in Section 1.

The following sub-sections describe multihoming examples to illustrate the PCP client behavior.

A.1. IPv6 Multihoming

In this example of an IPv6 multihomed network, two or more routers co-located with firewalls are present on a single link shared with the host(s). Each router is, in turn, connected to a different service provider network, and the host in this environment would be offered multiple prefixes and advertised multiple DNS servers. Consider a scenario in which firewalls within an IPv6 multihoming environment also implement a PCP server. The PCP client learns the available PCP servers using DHCP [RFC7291] or any other provisioning mechanism. In reference to Figure 2, a typical model is to embed DHCP servers in rtr1 and rtr2. A host located behind rtr1 and rtr2 can contact these two DHCP servers and retrieve from each server the IP address(es) of the corresponding PCP server.

The PCP client will send PCP requests in parallel to each of the PCP servers.

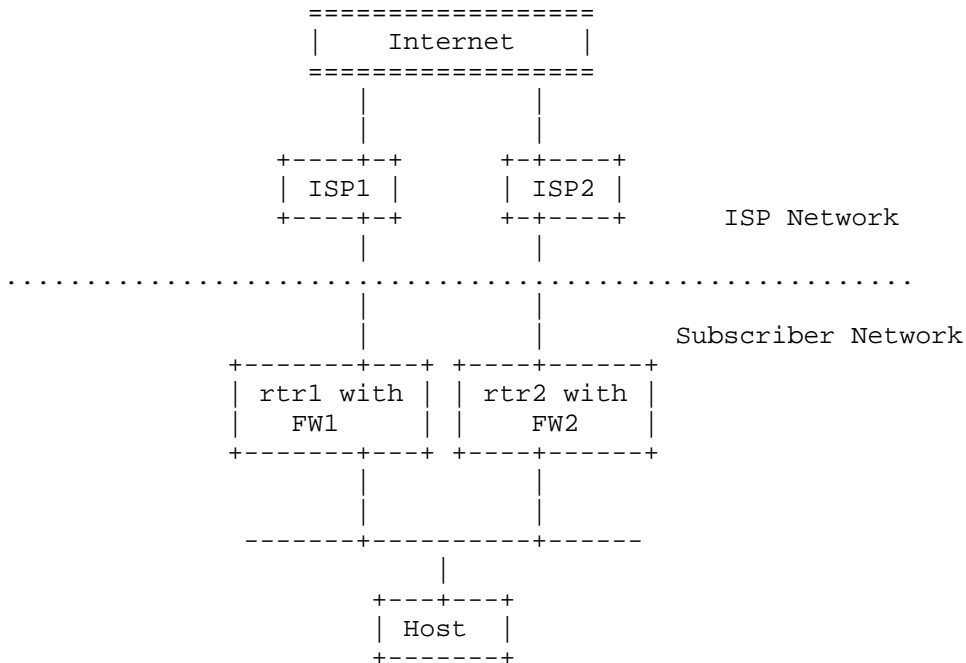


Figure 2: IPv6 Multihoming

A.2. IPv4 Multihoming

In this example of an IPv4 multihomed network described in "NAT- or RFC2260-based Multihoming" (Section 3.3 of [RFC4116]), the gateway router is connected to different service provider networks. This method uses Provider-Aggregatable (PA) addresses assigned by each transit provider to which the site is connected. The site uses NAT to translate the various provider addresses into a single set of private-use addresses within the site. In such a case, two PCP servers might have to be present to configure NAT to each of the transit providers. The PCP client learns the available PCP servers using DHCP [RFC7291] or any other provisioning mechanism. In reference to Figure 3, a typical model is to embed the DHCP server and the PCP servers in rtr1. A host located behind rtr1 can contact the DHCP server to obtain IP addresses of the PCP servers. The PCP client will send PCP requests in parallel to each of the PCP servers.

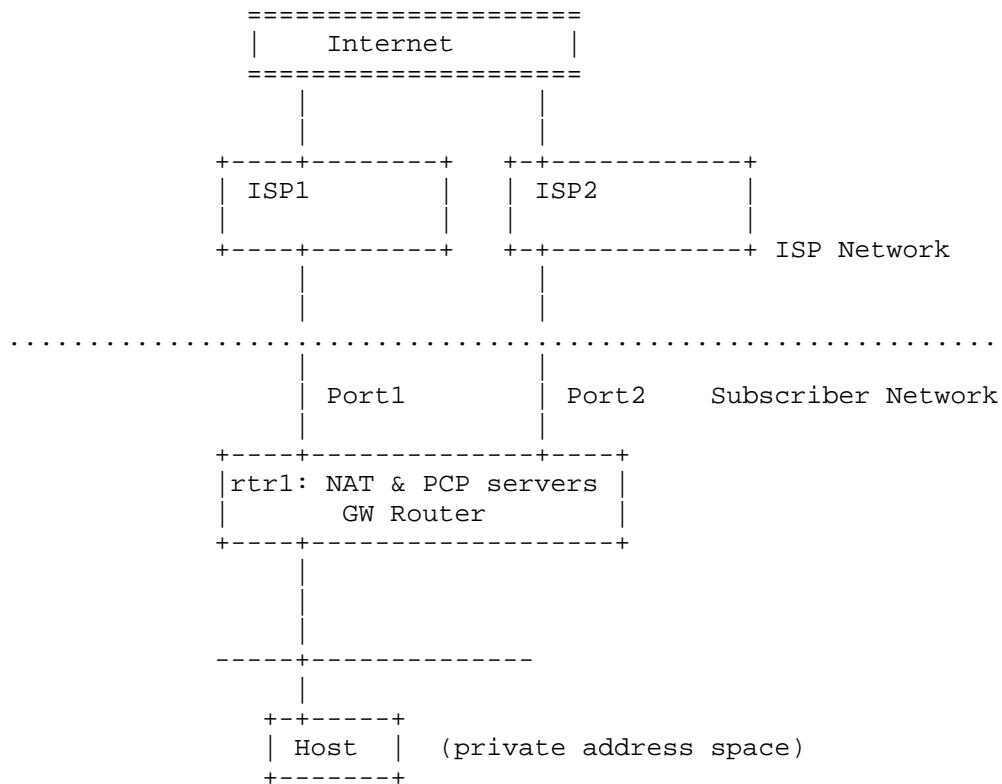


Figure 3: IPv4 Multihoming

Acknowledgements

Many thanks to Dave Thaler, Simon Perreault, Hassnaa Moustafa, Ted Lemon, Chris Inacio, and Brian Haberman for their reviews and comments.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes 35000
France

EEmail: mohamed.boucadair@orange.com

Reinaldo Penno
Cisco Systems, Inc.
United States

EEmail: repenno@cisco.com

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

EEmail: dwing@cisco.com

Prashanth Patil
Cisco Systems, Inc.
Bangalore
India

EEmail: praspati@cisco.com

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

EEmail: tireddy@cisco.com