PCP Working Group
Internet-Draft
Updates: 6887 (if approved)
Intended status: Standards Track
Expires: July 10, 2014

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

January 06, 2014

PCP Server Selection
draft-ietf-pcp-server-selection-02

Abstract

Multiple IP addresses may be configured on a PCP client in some deployment contexts such as multi-homing. This document specifies the behavior to be followed by the PCP client to contact its PCP server(s) when one or several PCP server IP addresses are configured.

This document updates RFC6887.

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

Table of Contents

1. Introduction

Multiple IP addresses may be configured on a PCP client in some deployment contexts such as multi-homing (see Appendix A). A PCP server may also have multiple IP addresses associated with it. This document specifies the behavior to be followed by the PCP client [RFC6887] to contact its PCP server(s) [RFC6887] when it receives one or several PCP server IP addresses (e.g., using DHCP [I-D.ietf-pcp-dhcp]).

This document is not specific to DHCP; any other mechanism can be used to configure PCP server IP addresses.

It is out of scope of this document to enumerate all deployment scenarios that require multiple IP addresses to be configured.

2. Terminology

This document makes use of the following terms:

- PCP server denotes a functional element that 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., Network Address Translation (NAT), firewall) it controls. Refer to [RFC6887].

- PCP client denotes a PCP software instance responsible for issuing PCP requests to a PCP server. Refer to [RFC6887].

3. IP Address Selection

These steps specify the behavior to be followed by the PCP client to contact a PCP server when the PCP client has multiple IP addresses for a single PCP server. Additional considerations to be taken into account when the PCP client is multi-interfaced are specified in Section 4:

1. If the PCP client can use both address families when communicating to a particular PCP server, the PCP client SHOULD select the source address of the PCP request to be of the same IP address family as its requested PCP mapping (i.e., the address family of the Requested External IP Address).

2. Whenever communicating with a PCP server, the rules of Section 6 of [RFC6724] SHOULD be followed by using the source address 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 message to the PCP server's IP address following the retransmission procedure specified in Section 8.1.1 of [RFC6887]. If no response is received after MRC attempts, the PCP client re-tries the procedure excluding the destination addresses which did not respond. The PCP client SHOULD ignore any response received from an IP address after exhausting MRC attempts for that particular IP address. If, when sending PCP requests, the PCP client receives a hard ICMP error [RFC1122] it SHOULD 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 is repeated every fifteen 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, it sends subsequent PCP requests to the same server's IP address until that IP address becomes non-responsive, which causes the PCP client to follow the steps above to contact its PCP server.

For efficiency, the PCP client should use the same Mapping Nonce for requests sent to all PCP server' IP addresses.

If several PCP servers are configured, each with multiple IP addresses, the PCP client contacts all PCP servers in parallel following the steps described above. 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 is implementation-specific and only the PCP client can decide whether all mappings are needed or only a subset of them.

4. Multiple Interfaces

When an end host has multiple interfaces concurrently active (e.g., IEEE 802.11 and 3G), a PCP client would discover different PCP servers over different interfaces. A host may have multiple network interfaces (e.g, 3G, IEEE 802.11, etc.); each configured differently. Each PCP server learned MUST be associated with the interface via which it was learned. Particularly, the PCP client relies on the source IP address of an outgoing PCP request to select which PCP server(s) to use.

Although multiple interfaces may be available, a PCP client might choose to use just one based on, for example, cost and bandwidth requirements, and therefore would need to send PCP requests to just one PCP server.

Generic multi-interfaces considerations are documented in Section 8.4 of [RFC6887]

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 Section 3.

```
                              ISP Network
                            |              |
       ........................................................
                            |              |        Subscriber Network
                    +-------+------+  +----+---------+
                    | PCP-Server-A |  | PCP-Server-B |
                    |              |  |              |
                    +-------+------+  +----+---------+
           192.0.2.1        |              |     198.51.100.1
           2001:db8:2222::1 |              |     2001:db8:3333::1
                            |              |
                            |              |
                     -------+--------------+-----------
                                    |
                                    |    203.0.113.0
                                    |    2001:db8:1111::1
                                +---+---+
                                | Host  |
                                +-------+
```

Figure 1

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

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

Suppose that:

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

o The path to reach 2001:db8:2222::1 is working

The PCP client 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 4 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 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 to the PCP client. It is the responsibility of any PCP server provisioning document(s) to elaborate on the security considerations to discover a legitimate PCP server.

7. IANA Considerations

This document does not request any action from IANA.

8. Acknowledgements

Many thanks to Dave Thaler for the review and comments.

9. References

9.1. Normative References

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

[RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.

[RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, April 2013.

9.2. Informative References

[I-D.ietf-pcp-dhcp]
Boucadair, M., Penno, R., and D. Wing, "DHCP Options for the Port Control Protocol (PCP)", draft-ietf-pcp-dhcp-09 (work in progress), November 2013.

[RFC1122] Braden, R., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.

[RFC1123] Braden, R., "Requirements for Internet Hosts - Application and Support", STD 3, RFC 1123, October 1989.

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

[RFC4116] Abley, J., Lindqvist, K., Davies, E., Black, B., and V. Gill, "IPv4 Multihoming Practices and Limitations", RFC 4116, July 2005.

Appendix A. Multi-homing

The main problem of a PCP multi-homing 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 in parallel with the following assumptions:

- o There is no requirement that multiple PCP servers have the same capabilities.
- o PCP requests to different servers are independent, the result of a PCP request to one server does not influence another.
- o If PCP servers provide NAT, it is out of scope how the client manages ports across PCP servers. For example, whether PCP client requires all external ports to be the same or whether there are ports available at all.

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

A.1. IPv6 Multi-homing

In this example of an IPv6 multi-homed 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 multi-homing environment also implement a PCP server. The PCP client learns the available PCP servers using DHCP [I-D.ietf-pcp-dhcp] 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 server and retrieves from each server the IP address of the corresponding PCP server.

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

```
                          ==================
                          |    Internet    |
                          ==================
                             |          |
                             |          |
                        +----+-+      +-+----+
                        | ISP1 |      | ISP2 |
                        +----+-+      +-+----+      ISP Network
                             |          |
       ..........................................................
                             |          |
                             |          |        Subscriber Network
                     +-------+---+ +----+------+
                     | rtr1 with | | rtr2 with |
                     |   FW1     | |    FW2    |
                     +-------+---+ +----+------+
                             |          |
                             |          |
                      -------+----------+------
                                  |
                              +---+---+
                              | Host  |
                              +-------+
```

Figure 2: IPv6 Multihoming

A.2. IPv4 Multi-homing

In this example an IPv4 multi-homed network described in 'NAT- or RFC2260-based multi-homing' (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 have to be present to control NAT to each of the transit providers. The PCP client learns the available PCP servers using DHCP [I-D.ietf-pcp-dhcp] 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.

```
                        =====================
                        |    Internet       |
                        =====================
                           |              |
                           |              |
                      +----+--------+   +-+------------+
                      | ISP1        |   | ISP2         |
                      |             |   |              |
                      +----+--------+   +-+------------+ ISP Network
                           |              |
                           |              |
         ...............................................................
                           |              |
                           | Port1        | Port2    Subscriber Network
                           |              |
                      +----+-------------------+
                      |rtr1: NAT & PCP servers |
                      |       GW Router        |
                      +----+-------------------+
                           |
                           |
                           |
                      -----+--------------
                           |
                         +-+-----+
                         | Host  |  (private address space)
                         +-------+
```

Figure 3: IPv4 Multihoming

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

Prashanth Patil
Cisco Systems, Inc.
Bangalore
India

EMail: praspati@cisco.com

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

EMail: tireddy@cisco.com