PCP Working Group
Internet-Draft
Updates: 6887 (if approved)
Intended status: Standards Track
Expires: November 23, 2013

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

May 22, 2013

PCP Server Selection
draft-ietf-pcp-server-selection-01

Abstract

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

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

Copyright Notice

Table of Contents

1. Introduction

This document specifies the behavior to be followed by the PCP client [RFC6887] to contact its PCP server(s) [RFC6887] when receiving one or several PCP server names (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 names.

Multiple names may be configured to a PCP client in some deployment contexts such as multi-homing (see Appendix A). It is out of scope of this document to enumerate all deployment scenarios which require multiple names to be configured.

This document assumes appropriate name resolution means (e.g., Section 6.1.1 of [RFC1123]) are available on the host client.

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., 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].
- Name is a string that can be passed to getaddrinfo (Section 6.1 of [RFC3493]), such as a DNS name, address literals, etc. 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).

3. Name Resolution

Each configured 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). Then, the PCP client follows the procedure specified in Section 4 to contact its PCP server(s).

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.

4. IP Address Selection

This section specifies the behavior to be followed by the PCP client to contact its PCP server(s) when receiving one or several PCP server names:

1. If only one PCP server name is configured: if a list of IP addresses is returned as a result of resolving the PCP server name, the PCP client follows the procedure specified in Section 4.1.

2. If several PCP server names are configured: each name is treated as a separate PCP server. Moreover, each name may be resolved into one IP address or a list of IP addresses. The PCP client contacts in parallel the first IP address of each name and follows the procedure specified in Section 4.1 for the list of IP addresses returned for each name. Section 5 provides some examples to illustrate this procedure.

This procedure does not require any knowledge of the capabilities of the PCP-controlled device(s). Instead, the PCP client assumes each configured name refer to a separate PCP server.

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 deployment-specific. Only the PCP client can decide whether all the mappings are needed or only a subset of them.

4.1. Serial Queries

The PCP client initializes its Maximum Retransmission Count (MRC) to 4.

The PCP client sends its PCP message to the PCP server following the retransmission procedure specified in Section 8.1.1 of [RFC6887]. If no response is received after MRC attempts, the PCP client tries with the next IP address in its list of PCP server addresses. If it has exhausted its list, the procedure is repeated every fifteen minutes until the PCP request is successfully answered. If, when sending PCP requests the PCP client receives an ICMP error (e.g., port unreachable, network unreachable) it SHOULD immediately try the next IP address in the list. Once the PCP client has successfully received a response from a PCP server address on that interface, it sends subsequent PCP requests to that same server address until that PCP server becomes non-responsive, which causes the PCP client to attempt to re-iterate the procedure starting with the first PCP server address on its list.

5. Examples

The following sub-sections provide three examples to illustrate the procedure.

For all these examples, let's suppose pcpserver-x, pcpserver-y and pcpserver-z are configured as PCP server names.

5.1. Example 1

Let's also suppose:

* IPx1 and IPx2 are returned for pcpserver-x; IPx1 is not reachable.
* IPy1 and IPy2 are returned for pcpserver-y; IPy1 is reachable
* IPz1 and IPz2 are returned for pcpserver-z; IPz1 is reachable

The procedure to contact the PCP servers is as follows:

* Send PCP requests to all servers: IPx1, IPy1 and IPz1
* Responses are received from IPy1 and IPz1 but not from IPx1
  - The request is re-sent to IPx1
  - If no response is received after four attempts, the request is sent to IPx2

5.2. Example 2

Now, if the following conditions are made:

* IPx1 and IPx2 are returned for pcpserver-x; IPx1 is not reachable.
* IPy1 and IPy2 are returned for pcpserver-y; IPy1 is reachable
* IPz1 and IPz2 are returned for pcpserver-z; IPz1 is not reachable

The procedure to contact the PCP servers lead to the following:

* Send PCP requests to all servers: IPx1, IPy1 and IPz1
* A response is received from IPy1 but not from IPx1 and IPz1
  - the requests are re-sent to IPx1 and IPz1
  - If no response is received after four attempts, the request is then sent to IPx2 and IPz2

5.3. Example 3

Let's suppose now that:

* IPx1 and IPx2 are returned for pcpserver-x; IPx1 is not reachable.
* IPy1 and IPy2 are returned for pcpserver-y; IPy1 is not reachable
* IPz1 and IPz2 are returned for pcpserver-z; IPz1 is not reachable

The procedure to contact the PCP servers is as follows:

* Send PCP requests to all servers: IPx1, IPy1 and IPz1
* No answer is received for all requests
  - the requests are re-sent to IPx1, IPy1 and IPz1

- If no response is received after four attempts, the request is then sent to IPx2, IPy2 and IPz2

6. Security Considerations

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

This document does not specify how PCP server names are provisioned to the PCP client. It is the responsibility of 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 D. 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.

[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-07 (work in progress), March 2013.

[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.

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

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 4, if a PCP client discovers multiple PCP server names, 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 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/NTP servers. Consider a scenario in which firewalls within an IPv6 multi-homing environment also implement a PCP server. PCP client learns of the available PCP servers using DHCP [I-D.ietf-pcp-dhcp] or any other PCP server discovery technique defined in future specifications. 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    |
                  +-------+---+ +----+------+
                          |          |
                          |          |
                          |          |
                   -----+-+-----+------
                          |
                        +-+-----+
                        | Hosts |
                        +-------+
```

Figure 1: 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 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. PCP client learns of the available PCP servers using DHCP [I-D.ietf-pcp-dhcp] or any other PCP server discovery technique defined in future specifications. The PCP client will send PCP requests in parallel to each of the PCP servers.

```
                     ====================
                     |    Internet       |
                     =====================
                        |              |
                        |              |
                   +----+--------+   +-+------------+
                   | ISP1        |   | ISP2         |
                   |             |   |              |
                   +----+--------+   +-+------------+ ISP Network
                        |              |
                        |              |
```

```
      .............................................................
                        |              |
                        | Port1        | Port2    Subscriber Network
                        |              |
                   +----+----------------+
                   |   NAT & PCP servers  |
                   |       GW Router      |
                   +----+----------------+
                        |
                        |
                        |
                   -----+-+-----+------
                        |
                      +-+-----+
                      | Hosts |  (private address space)
                      +-------+
```

Figure 2: IPv4 Multihoming

A.3. Multiple interfaces and Servers

In case for multi-homing when an end host such as a mobile terminal has multiple interfaces concurrently active (e.g., IEEE 802.11 and 3G), a PCP client would discover different PCP servers over different interfaces. Although multiple interfaces are available, an application 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.

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