PCP Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 18, 2013

M. Boucadair
France Telecom
R. Dupont
Internet Systems Consortium
R. Penno
D. Wing
Cisco

August 17, 2012

Port Control Protocol (PCP) Proxy Function
draft-ietf-pcp-proxy-01

Abstract

This document specifies a new PCP functional element denoted as PCP Proxy. The PCP Proxy relays PCP requests received from PCP Clients to upstream PCP Server(s). This function is mandatory when PCP Clients can not be configured with the address of the PCP Server located more than one hop.

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 February 18, 2013.

Table of Contents

1. Introduction

This document defines a new PCP [I-D.ietf-pcp-base] function element, called PCP Proxy, which is meant to facilitate the communication between a PCP Client and upstream PCP Server(s). The PCP Proxy acts as a PCP Server receiving PCP requests on internal interfaces, and as a PCP Client forwarding accepted PCP requests on an external interface to a PCP Server. The PCP Server in turn send PCP responses to the PCP Proxy external interface which are finally forwarded to PCP Clients. A reference architecture is depicted in Figure 1.

A PCP Proxy can be for instance embedded in a CP (Customer Premises) router while the PCP Server is located in a network operated by an ISP (Internet Service Provider). It is out of scope of this document to list all deployment scenarios requiring a PCP Proxy to be involved.

The PCP Proxy can be simple, i.e., implement as transparent/minimal processing as possible, or it can support advanced features (see Section 10). A Proxy can be co-located with UPnP IGD [I-D.ietf-pcp-upnp-igd-interworking] or/and NAT-PMP [I-D.bpw-pcp-nat-pmp-interworking] Interworking Function (IWF).

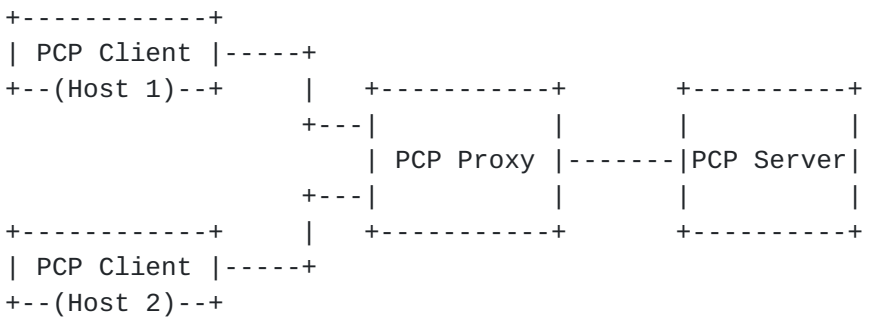

Figure 1: Reference Architecture

2. 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].

3. PCP Server Discovery and Provisioning

The PCP Proxy MUST implement one of the discovery methods listed in

[I-D.ietf-pcp-base] (e.g., DHCP [I-D.ietf-pcp-dhcp]).

The address of the PCP Proxy is provisioned to local PCP Clients as their default PCP Server: If the PCP DHCP option [I-D.ietf-pcp-dhcp] is supported by an internal PCP Client, it will retrieve the PCP Server IP address to use from its local DHCP server; otherwise internal PCP Clients will assume their default router being the PCP Server.

4. PCP Proxy as a PCP Server

The PCP Proxy acts as a PCP Server for internal hosts and accepts PCP requests on the interface(s) facing them, e.g., it creates servicing socket(s) and bound them to each address of this (these) interface(s) on UDP port 5350.

When the topology makes a routing loop possible, the PCP Proxy MAY check it is not the source of a PCP message it received.

5. Control of the Firewall

A security policy to accept PCP messages from the provisioned PCP Server(s) is to be enabled on the device embedding the PCP Proxy. This policy can be for instance triggered by DHCP configuration or by outbound PCP requests issued from the PCP Proxy to the provisioned PCP Server.

In order to accept inbound and outbound traffic associated with PCP mappings instantiated in the upstream PCP Server, appropriate security policies are to be configured on the firewall.

For instance if the firewall rules have a lifetime, PCP response can be snooped in order to instantiate the corresponding firewall rules with the same lifetime. If they have no lifetime, an explicit dynamic mapping table can be kept in the PCP Proxy state in order to instantiate and remove corresponding firewall rules.

REMOTE_PEER_FILTER Options can be installed into the local firewall, forwarded to the PCP Server so installed into the remote NAT/firewall or both.

6. No NAT is Co-located with the PCP Proxy

When no NAT is co-located with the PCP Proxy, the port numbers included in received PCP messages (from the PCP Server or PCP

Client(s)) are not altered by the PCP Proxy. Nevertheless, the PCP Client IP Address MUST be changed to the address of the PCP Proxy and a THIRD_PARTY Option inserted to carry the IP address of the source PCP Client.

Because no NAT is invoked, there is no reachability failure risk to relay to the PCP Server unknown Options and OpCodes which carry an IP address.

7. PCP Proxy Co-located with a NAT Function

When the PCP Proxy is co-located with a NAT function, it MUST update the content of received requests with the mapped port number and the address belonging to the external interface of the PCP Proxy (i.e., after the NAT operation) and not as initially positioned by the PCP Client. For the reverse path, PCP responses MUST be updated by the PCP Proxy to replace the internal port number to what has been initially positioned by the PCP Client. For this purpose the PCP Proxy MUST have an access to the local NAT state maintained locally. Because PCP messages with an unknown OpCode or Option can carry a hidden internal address or internal port which will not be translated:

- a PCP Proxy co-located with a NAT SHOULD reject by an UNSUPP_OPCODE error response a received request with an unknown OpCode;

- a PCP Proxy co-located with a NAT SHOULD reject by an UNSUPP_OPTION error response a received request with a mandatory-to-process unknown Option;

- a PCP Proxy co-located with a NAT MAY remove any optional-to-process unknown Options from received requests before forwarding them.

Rejecting unknown Options and OpCodes has the drawback of preventing a PCP Client to make use of new capabilities offered by the PCP Server but not supported by the PCP Proxy even if no IP address and/or port is included in the Option/OpCode.

When a PCP request is received and accepted by the PCP Proxy the corresponding mapping (explicit dynamic mapping for a MAP request, implicit dynamic mapping for a PEER request) is looked for in the local NAT state and temporary created if it does not exist. Temporary means it is deleted if no SUCCESS response is received, either explicitly or because of its short lifetime at creation.

If the local NAT associates explicit dynamic mappings to a lifetime, the requested lifetime in MAP requests SHOULD be adjusted to be in the accepted range of the local NAT, and the assigned lifetime copied from MAP responses to the corresponding mapping in the local NAT. The same processing applies to implicit dynamic mappings and PEER requests/responses.

Otherwise explicit dynamic mappings have an undefined lifetime in the local NAT and the PCP Proxy SHOULD maintain an explicit dynamic mapping table and SHOULD delete corresponding explicit dynamic mappings in the local NAT when they expire or are deleted by the MAP request with a zero requested lifetime.

8. MAP/PEER Handling

A simple PCP Proxy performs minimal modifications to PCP requests and responses, in particular it does not change the Nonce value in requests and the Epoch value in responses. A simple PCP Proxy is assumed to handle only one PCP Server.

The detailed behavior at the reception of a PCP request on an internal interface is as follows:

- Check if the source IP address and the PCP Client IP Address are the same.
- Apply security controls, including with the result of the previous item.
- If the request is rejected, build a synthetic error response and send it back to the PCP Client.
- If the request is accepted, adjust it (e.g., adding a THIRD_PARTY Option, updating the PCP Client IP Address and Internal Port to their translated values as specified in Section 7 and forward it on a fresh UDP socket connected to the PCP Server).
- Wait for the response during a reasonable delay.
- When the response is received from the PCP Server, adjust it back (e.g., removing the THIRD_PARTY Option added previously, updating the PCP Client IP Address and Internal Port to their initial values as specified in Section 7), forward it to the source PCP Client and close the socket to the PCP Server.
- On a hard error on the UDP socket, build a synthetic ICMP error and send it to the source PCP Client.

The reasonable delay minimum value is 20 seconds, request retransmission is handled by PCP clients.

For each pending request, the proxy MUST maintain in a data record:

- the request payload
- the interface where the request was received
- the source IP address of the request
- the source UDP port of the request
- the UDP socket connected to the PCP server
- an expire timeout

Receiving interfaces can be implemented by a set of servicing sockets, each socket bound to an address of an internal interface. Interface, source address and port are used to send back packets to the source PCP Client. The request payload is used to generate synthetic ICMP. Responses are received on the UDP socket.

Too large requests SHOULD be forwarded to the PCP Server in order to relay back the error response, i.e., the PCP Proxy is not in charge to enforce the message size limit and in general the PCP Proxy SHOULD NOT generate error response for a reason other than security controls. No behavior is specified in the case the PCP Proxy processing (e.g., adding a THIRD_PARTY Option) makes a valid request too large when it is sent to the PCP Server.

9. Mapping Repair

ANNOUNCE requests received from PCP Clients are handled locally; as such these requests MUST NOT be relayed to the provisioned PCP Server.

Upon receipt of an unsolicited ANNOUNCE response from a PCP Server, the PCP Proxy proceeds to renewing the mappings and checks whether there are changes compared to a local cache if it is maintained by the PCP Proxy. If no change is detected, no unsolicited ANNOUNCE is generated towards PCP Clients. If a change is detected, the PCP Proxy MUST generate unsolicited ANNOUNCE message(s) to appropriate PCP Clients. If the PCP Proxy does not maintain a local cache for the mappings, unsolicited ANNOUNCE messages are relayed to PCP Clients.

Unsolicited PCP MAP/PEER responses received from a PCP Server are handled as any normal MAP/PEER response. To handle unsolicited PCP MAP/PEER responses, the PCP Proxy is required to maintain a local cache of instantiated mappings in the PCP Server (Section 10.5).

Upon change of its external IP address, the PCP Proxy SHOULD renew the mappings it maintained. If the PCP Server assigns a different external port, the PCP Proxy SHOULD follow the mapping repair procedure defined in [I-D.ietf-pcp-base]. This can be achieved only if a full state table is maintained by the PCP Proxy.

10. Advanced Functions

Below are listed a set of advanced features which may be supported by the PCP Proxy.

10.1. Multiple PCP Servers

A PCP Proxy MAY offer to handle multiple PCP Servers at the same time, each PCP Server is associated to each own handled Epoch value according to Section 10.2. PCP Clients are not aware of the presence of multiple PCP Servers.

According to [I-D.ietf-pcp-dhcp], if several PCP Names are configured to the PCP Proxy, it will contact in parallel all these PCP Servers.

In some contexts (e.g., PCP-controlled CGNs), the PCP Proxy MAY load balance the PCP Client among available PCP Servers. The PCP Proxy MUST ensure requests of a given PCP Client are relayed to the same PCP Server.

In other deployment scenarios (e.g., presence of multiple PCP-controlled firewalls), the PCP Proxy MUST relay PCP requests to all these PCP Servers.

10.2. Epoch Handling

A PCP Proxy MAY use its own internal timers and not blindly copy them from PCP responses. There should be no advantages to have more than one managed Epoch per PCP Server.

The Epoch MUST be reset when explicit dynamic mappings are lost, i.e.:

- at startup if the PCP Proxy can't recover the state.

- when the WAN address is changed or any similar events which show any previous state is no longer valid.

- when the Epoch value in a PCP response is too small (cf. Epoch value validation rules in [I-D.ietf-pcp-base]).

- when the External IP Address has changed.

The last two rules are per PCP Server, a PCP Proxy MAY check these conditions in all received responses for a PCP Server.

10.3. Request/Response Caching

A PCP Proxy providing request/response caching checks each time it receives a PCP request if it has already seen the same request recently and got the corresponding PCP response. In this case, it sends back directly the cached response with the proper Epoch value and not forward the request to the PCP Server.

10.4. Retransmission Handling

An extension of the previous service is to manage the retransmission of pending requests to the PCP Server internally, i.e., no longer driven by the PCP Client. A cache entry SHOULD be expired after a delay short enough to keep it easy to distinguish it from a replay.

10.5. Full State

A PCP Proxy MAY keep the full state, i.e., an image of all active explicit dynamic mappings is kept in memory. When this service is supported the state SHOULD be recovered in case of failures (e.g., according to [I-D.boucadair-pcp-failure]).

11. IANA Considerations

This document makes no request of IANA.

12. Security Considerations

The PCP Proxy MUST follow the security considerations elaborated in [I-D.ietf-pcp-base] for both the client and server side.

A received request carrying an unknown OpCode or Option SHOULD be dropped (or in the case of an unknown Option which is not mandatory-to-process the Option be removed) if it is not a priori compatible with security controls or correct processing.

13. References

13.1. 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-26 (work in progress), June 2012.

[I-D.ietf-pcp-dhcp]
Boucadair, M., Penno, R., and D. Wing, "DHCP Options for the Port Control Protocol (PCP)", draft-ietf-pcp-dhcp-04 (work in progress), August 2012.

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

13.2. Informative References

[I-D.boucadair-pcp-failure]
Boucadair, M., Dupont, F., and R. Penno, "Port Control Protocol (PCP) Failure Scenarios", draft-boucadair-pcp-failure-03 (work in progress), May 2012.

[I-D.bpw-pcp-nat-pmp-interworking]
Boucadair, M., Penno, R., Wing, D., and F. Dupont, "Port Control Protocol (PCP) NAT-PMP Interworking Function", draft-bpw-pcp-nat-pmp-interworking-00 (work in progress), March 2011.

[I-D.ietf-pcp-upnp-igd-interworking]
Boucadair, M., Dupont, F., Penno, R., and D. Wing, "Universal Plug and Play (UPnP) Internet Gateway Device (IGD)-Port Control Protocol (PCP) Interworking Function", draft-ietf-pcp-upnp-igd-interworking-02 (work in progress), August 2012.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Francis Dupont
Internet Systems Consortium

Email: fdupont@isc.org

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