

PCP working group
Internet-Draft
Intended status: Standards Track
Expires: November 21, 2015

S. Kiesel
University of Stuttgart
R. Penno
Cisco Systems, Inc.
S. Cheshire
Apple
May 20, 2015

Port Control Protocol (PCP) Anycast Addresses
draft-ietf-pcp-anycast-06

Abstract

The Port Control Protocol (PCP) Anycast Addresses enable PCP clients to transmit signaling messages to their closest PCP-aware on-path NAT, Firewall, or other middlebox, without having to learn the IP address of that middlebox via some external channel. This document establishes one well-known IPv4 address and one well-known IPv6 address to be used as PCP Anycast Addresses.

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 November 21, 2015.

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](#) [3](#)
- [2. PCP Server Discovery based on well-known IP Address](#) [4](#)
 - [2.1. PCP Discovery Client behavior](#) [4](#)
 - [2.2. PCP Discovery Server behavior](#) [4](#)
- [3. Deployment Considerations](#) [5](#)
- [4. IANA Considerations](#) [6](#)
 - [4.1. Registration of IPv4 Special Purpose Address](#) [6](#)
 - [4.2. Registration of IPv6 Special Purpose Address](#) [6](#)
- [5. Security Considerations](#) [7](#)
 - [5.1. Information Leakage through Anycast](#) [7](#)
 - [5.2. Hijacking of PCP Messages sent to Anycast Addresses](#) [7](#)
- [6. Acknowledgments](#) [8](#)
- [7. References](#) [9](#)
 - [7.1. Normative References](#) [9](#)
 - [7.2. Informative References](#) [9](#)
- [Authors' Addresses](#) [10](#)

1. Introduction

The Port Control Protocol (PCP) [[RFC6887](#)] provides a mechanism to control how incoming packets are forwarded by upstream devices such as Network Address Translator IPv6/IPv4 (NAT64), Network Address Translator IPv4/IPv4 (NAT44), and IPv6 and IPv4 firewall devices. Furthermore, it provides a mechanism to reduce application keep alive traffic [[I-D.ietf-pcp-optimize-keepalives](#)]. The PCP base protocol document [[RFC6887](#)] specifies the message formats used, but the address to which a client sends its request is either assumed to be the default router (which is appropriate in a typical single-link residential network) or has to be configured otherwise via some external mechanism, such as a configuration file or a DHCP option [[RFC7291](#)].

This document follows a different approach: it establishes two well-known anycast addresses for the PCP Server, one IPv4 address and one IPv6 address. These well-known addresses may be hard-coded into PCP clients. PCP clients usually send PCP requests to these addresses if no other PCP server addresses are known or after communication attempts to such other addresses have failed.

Using an anycast address is particularly useful in larger network topologies. For example, if the PCP-enabled NAT/firewall function is not located on the client's default gateway, but further upstream in a Carrier-grade NAT (CGN), sending PCP requests to the default gateway's IP address will not have the desired effect. When using a configuration file or the DHCP option to learn the PCP server's IP address, this file or the DHCP server configuration must reflect the network topology, and the router and CGN configuration. This may be cumbersome to achieve and maintain. If there is more than one upstream CGN and traffic is routed using a dynamic routing protocol such as OSPF, this approach may not be feasible at all, as it cannot provide timely information on which CGN to interact with. In contrast, when using the PCP anycast address, the PCP request will travel through the network like any other packet, without any special support from DNS, DHCP, other routers, or anything else, until it reaches the PCP-capable device, which receives it, handles it, and sends back a reply. A further advantage of using an anycast address instead of a DHCP option is, that the anycast address can be hard-coded into the application. There is no need for an application programming interface for passing the PCP server's address from the operating system's DHCP client to the application. For further discussion of deployment considerations see [Section 3](#).

2. PCP Server Discovery based on well-known IP Address

2.1. PCP Discovery Client behavior

The PCP anycast addresses, as defined in Sections [4.1](#) and [4.2](#), are added after the default router list (for IPv4 and IPv6) to the list of PCP server(s) (see [Section 8.1](#), step 2. of [[RFC6887](#)]). This list is processed as specified in [[RFC7488](#)].

Note: If, in some specific scenario, it was desirable to use only the anycast address (and not the default router), this could be achieved by putting the anycast address into the configuration file, or DHCP option, etc.

2.2. PCP Discovery Server behavior

A PCP Server can be configured to listen on the anycast address for incoming PCP requests.

PCP responses are sent from that same IANA-assigned address (see Page 6 of [[RFC1546](#)] for further discussion).

3. Deployment Considerations

There are known limitations when there is more than one PCP-capable NAT/firewall in a cascaded alignment, or in a parallel layout with asymmetric routing, or similar scenarios. Mechanisms to deal with those situations, such as state synchronization between PCP servers, are beyond the scope of this document.

For general recommendations regarding operation of anycast services see [[RFC4786](#)].

4. IANA Considerations

4.1. Registration of IPv4 Special Purpose Address

IANA is requested to assign a single IPv4 address from the 192.0.0.0/24 prefix and register it in the IANA IPv4 Special-Purpose Address Registry [[RFC6890](#)].

Attribute	Value
Address Block	192.0.0.???/32 (??? = TBD by IANA)
Name	Port Control Protocol Anycast
RFC	This document, if approved (TBD)
Allocation Date	Date of approval of this document (TBD)
Termination Date	N/A
Source	True
Destination	True
Forwardable	True
Global	True
Reserved-by-Protocol	False

4.2. Registration of IPv6 Special Purpose Address

IANA is requested to assign a single IPv6 address from the 2001:0000::/23 prefix and register it in the IANA IPv6 Special-Purpose Address Registry [[RFC6890](#)].

Attribute	Value
Address Block	2001:0????????/128 (??? = TBD by IANA)
Name	Port Control Protocol Anycast
RFC	This document, if approved (TBD)
Allocation Date	Date of approval of this document (TBD)
Termination Date	N/A
Source	True
Destination	True
Forwardable	True
Global	True
Reserved-by-Protocol	False

5. Security Considerations

In addition to the security considerations in [[RFC6887](#)], two additional issues are considered here.

5.1. Information Leakage through Anycast

In a network without any border gateway, NAT or firewall that is aware of the PCP anycast address, outgoing PCP requests could leak out onto the external Internet, possibly revealing information about internal devices.

Using an IANA-assigned well-known PCP anycast address enables border gateways to block such outgoing packets. In the default-free zone, routers should be configured to drop such packets. Such configuration can occur naturally via BGP messages advertising that no route exists to said address.

Sensitive clients that do not wish to leak information about their presence can set an IP TTL on their PCP requests that limits how far they can travel into the public Internet.

5.2. Hijacking of PCP Messages sent to Anycast Addresses

The anycast addresses are treated by normal host operating systems just as normal unicast addresses, i.e., packets destined for an anycast address are sent to the default router for processing and forwarding. Hijacking such packets in the first network segment would effectively require to impersonate the default router, e.g., by means of ARP spoofing in an Ethernet network. If such attacks are a serious concern in a given scenario, much more severe consequences to other protocols have to be feared as well. Therefore, adequate measures have to be taken to prevent spoofing attacks targeted at the default router.

Once an anycast message is forwarded closer to the core network, routing will likely become subject to dynamic routing protocols such as OSPF or BGP. Anycast messages could be hijacked by announcing counterfeited messages in these routing protocols. But again, an attacker capable of performing these attacks could cause significantly more damage to other protocols and therefore adequate means should be taken to prevent these attacks.

In addition to following best current practices in first hop security and routing protocol security, PCP authentication [[I-D.ietf-pcp-authentication](#)] may be useful in some scenarios, although it might thwart the goal of fully automatic configuration in other scenarios.

6. Acknowledgments

The authors would like to thank the members of the PCP working group for contributions and feedback, in particular Mohamed Boucadair, Charles Eckel, Simon Perreault, Tirumaleswar Reddy, Markus Stenberg, Dave Thaler, and Dan Wing.

[7.](#) References

[7.1.](#) Normative References

- [RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", [RFC 6887](#), April 2013.
- [RFC6890] Cotton, M., Vegoda, L., Bonica, R., and B. Haberman, "Special-Purpose IP Address Registries", [BCP 153](#), [RFC 6890](#), April 2013.
- [RFC7488] Boucadair, M., Penno, R., Wing, D., Patil, P., and T. Reddy, "Port Control Protocol (PCP) Server Selection", [RFC 7488](#), March 2015.

[7.2.](#) Informative References

- [I-D.ietf-pcp-authentication]
Wasserman, M., Hartman, S., Zhang, D., and T. Reddy, "Port Control Protocol (PCP) Authentication Mechanism", [draft-ietf-pcp-authentication-07](#) (work in progress), December 2014.
- [I-D.ietf-pcp-optimize-keepalives]
Reddy, T., Patil, P., Isomaki, M., and D. Wing, "Optimizing NAT and Firewall Keepalives Using Port Control Protocol (PCP)", [draft-ietf-pcp-optimize-keepalives-06](#) (work in progress), May 2015.
- [RFC1546] Partridge, C., Mendez, T., and W. Milliken, "Host Anycasting Service", [RFC 1546](#), November 1993.
- [RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", [BCP 126](#), [RFC 4786](#), December 2006.
- [RFC7291] Boucadair, M., Penno, R., and D. Wing, "DHCP Options for the Port Control Protocol (PCP)", [RFC 7291](#), July 2014.

Authors' Addresses

Sebastian Kiesel
University of Stuttgart Information Center
Networks and Communication Systems Department
Allmandring 30
Stuttgart 70550
Germany

Email: ietf-pcp@skiesel.de

Reinaldo Penno
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, California 95134
USA

Email: repenno@cisco.com

Stuart Cheshire
Apple Inc.
1 Infinite Loop
Cupertino, California 95014
USA

Phone: +1 408 974 3207

Email: cheshire@apple.com

