PCP working group
Internet-Draft
Intended status: Standards Track
Expires: April 8, 2016

S. Kiesel
University of Stuttgart
R. Penno
Cisco Systems, Inc.
S. Cheshire
Apple
October 6, 2015

Port Control Protocol (PCP) Anycast Addresses
draft-ietf-pcp-anycast-08

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 April 8, 2016.

Copyright Notice

Table of Contents

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. PCP clients usually send PCP requests to these well-known addresses if no other PCP server addresses are known or after communication attempts to such other addresses have failed. The anycast addresses are allocated from pools of special-purpose IP addresses (see Section 4), in accordance with Section 3.4 of [RFC4085]. Yet, a means to disable or override these well-known addresses (e. g., a configuration file option) should be available in implementations.

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

PCP clients can add the PCP anycast addresses, which are defined in Sections 4.1 and 4.2, 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

PCP Servers can be configured to listen on the anycast addresses for incoming PCP requests. When a PCP server receives a PCP requests destined for an anycast address it supports, it sends the corresponding PCP replies using that same anycast address as the source address (see Page 6 of [RFC1546] for further discussion).

3. Deployment Considerations

For general recommendations regarding operation of anycast services see [RFC4786]. Architectural considerations of IP anycast are discussed in [RFC7094].

In some deployment scenarios, using PCP anycasting may have certain limitations, which can be overcome by using additional mechanisms or by using other PCP server discovery methods instead, such as DHCP [RFC7291] or a configuration file.

One important example is a network topology, in which a network is connected to one or more upstream network(s) via several parallel firewalls, each individually controlled by its own PCP server. Even if all of these PCP servers are configured for anycasting, only one will receive the messages sent by a given client, depending on the state of the routing tables.

As long as routing is always symmetric, i.e., all upstream and downstream packets from/to that client are routed through this very same firewall, communication will be possible as expected. If there is a routing change, a PCP client using PCP anycasting might start interacting with a different PCP server. From the PCP client's point of view this would be the same as a PCP server reboot and the client could detect it by examining the Epoch field during the next PCP response or ANNOUNCE message. The client would re-establish the firewall rules and packet flows could resume.

If, however, routing is asymmetric, upstream packets from a client traverse a different firewall than the downstream packets to that client. Establishing policy rules in only one of these two firewalls by means of PCP anycasting will not have the desired result of allowing bi-directional connectivity. One solution approach to overcome this problem is an implementation-specific mechanism to synchronize state between all firewalls at the border of a network, i.e., a PEER message sent to any of these PCP servers would establish rules in all firewalls. Another approach would be to use a different discovery mechanism (e.g., DHCP or a configuration file) that allows a PCP client to acquire a list of all PCP servers controlling the parallel firewalls and configure each of them individually.

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], [RFC4786], and [RFC7094], two further security 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 towards the public Internet. However, methods for choosing an appropriate TTL value, e.g., based on the assumed radius of the trusted network domain, is beyond the scope of this document.

Before sending PCP requests with possibly privacy-sensitive parameters (e.g., IP addresses and port numbers) to the PCP anycast addresses, PCP clients can send an ANNOUNCE request (without parameters; see Section 14.1 of [RFC6887]), in order to probe whether a PCP server consumes and processes PCP requests sent to that anycast address.

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 the attacker to impersonate the default router, e.g., by means of ARP spoofing in an Ethernet network. 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. When analyzing the risk and possible consequences of such attacks in a given network scenario, the probable impacts on PCP signaling need to be put into proportion with probable impacts on other protocols such as the actual application protocols.

In addition to following best current practices in first hop security and routing protocol security, PCP authentication [RFC7652] may be useful in some scenarios. However, the effort needed for a proper setup of this authentication mechanism (e.g., installing the right shared secrets or cryptograpic keys on all involved systems) may thwart the goal of fully automatic configuration by using PCP anycast. Therefore, this approach may be less suitable for scenarios with high trust between the operator of the PCP-controlled middlebox and all users (e.g., a residential gateway used only by family members) or if there is anyway rather limited trust that the middlebox will behave correctly (e.g., the Wifi in an airport lounge). In contrast, this scheme may be highly useful in scenarios with many users and a trusted network operator, such as a large corporate network or a university campus network, which uses several parallel NATs or firewalls to connect to the Internet. Therefore, a thorough analysis of the benefits and costs of using PCP authentication in a given network scenario is recommended.

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-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.

[RFC4085] Plonka, D., "Embedding Globally-Routable Internet Addresses Considered Harmful", BCP 105, RFC 4085, DOI 10.17487/RFC4085, June 2005, <http://www.rfc-editor.org/info/rfc4085>.

[RFC4786] Abley, J. and K. Lindqvist, "Operation of Anycast Services", BCP 126, RFC 4786, December 2006.

[RFC7094] McPherson, D., Oran, D., Thaler, D., and E. Osterweil, "Architectural Considerations of IP Anycast", RFC 7094, DOI 10.17487/RFC7094, January 2014, <http://www.rfc-editor.org/info/rfc7094>.

[RFC7291] Boucadair, M., Penno, R., and D. Wing, "DHCP Options for the Port Control Protocol (PCP)", RFC 7291, July 2014.

[RFC7652] Cullen, M., Hartman, S., Zhang, D., and T. Reddy, "Port Control Protocol (PCP) Authentication Mechanism", RFC 7652, DOI 10.17487/RFC7652, September 2015, <http://www.rfc-editor.org/info/rfc7652>.

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