

Port Control Protocol  
Internet-Draft  
Intended status: Standards Track  
Expires: April 22, 2012

R. Penno  
Juniper Networks  
October 20, 2011

PCP Support for Multi-Zone Environments  
draft-penno-pcp-zones-01

Abstract

A zone is a notion which denotes a routing instance, a set interfaces or prefixes characterized by having a different address realm and/or security policy. A NAT device can route packets with the same source IP address to different zones depending on configuration policies such as destination IP address. This functionality has been present for many years in NAT devices from multiple vendors. PCP allows a host to interact with a PCP-controlled NAT device and request an external IP and port. Therefore a PCP Server that controls the NAT device and receives a PCP request from a host needs to know from which NAT pool to allocate an external IP address and port. This document specifies an extension to PCP to support the zone concept.

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

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 22, 2012.

Copyright Notice

Copyright (c) 2011 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
1.1. Terminology . . . . .	3
1.2. Problem Statement . . . . .	3
1.3. Scope . . . . .	4
2. PCP Base Support for Multiple Zones . . . . .	4
2.1. PCP PEER Request . . . . .	4
2.2. PCP MAP Request . . . . .	5
3. PCP Extension for Multiple Zones . . . . .	5
4. IANA Considerations . . . . .	6
5. Security Considerations . . . . .	6
6. Acknowledgements . . . . .	6
7. References . . . . .	6
7.1. Normative References . . . . .	6
7.2. Informative References . . . . .	7
Author's Address . . . . .	8

## 1. Introduction

A zone is a routing instance, set interfaces or prefixes characterized by having a different address domain or security policy. A NAT device is present on each zone through NAT pools which are used to translate packet to and from a zone. The PCP protocol allows a host to interact with a NAT device and request a external IP and port. Since a NAT Device can route packets with the same source IP address to different Zones depending on policy or packet match conditions, the PCP Server that interacts with the NAT device and receives a PCP request from a host needs to know from which NAT pool to allocate an IP address and port.


### 1.1. Terminology

This document uses PCP terminology defined in [I-D.ietf-pcp-base]]. In addition the following terms are defined in this document:

- o Zone: A routing instance, set of interfaces or network prefixes that has a separate addressing domain or security policy.
- o Address Domain: A collection of IP addresses. A NAT device is present on each domain through one or more NAT pools associated with each Zone.

### 1.2. Problem Statement

A PCP Server can control a NAT attached to distinct zones; each zone is characterised by one or several address pools. In such environment the NAT must rely on a pre-configured policy to determine which address pool to use when handling an IP packet coming from an internal host. An example of such policy may be to rely on the destination IP address, DSCP value(s), protocol (e.g., SIP, RTP, RTSP), etc.



The core of the problem is that packets from the same source IP address can be routed to any of the zones depending on match conditions based on the 5-tuple. Moreover, sessions could be initiated from any of these zones toward the host. These zones many times have different addressing domains and therefore different NAT pools. This means that packets from the host will use a different NAT pool depending on the destination zone.

It is important to notice that zones (or similar concept) has been present in Enterprise NAT and CGN from multiple vendors for many years. It is the advent and interaction with PCP that has created a need for a standardized approach.

### 1.3. Scope

The matching conditions that ultimately decide where to route a packet can be very elaborate including even application layer information. But the scope of this document is to abstract such implementation specific approaches behind the concept of a Zone-ID.

## 2. PCP Base Support for Multiple Zones

Before discussing extensions to the PCP protocol in the following sections we discuss how to support multiple zones with the current methods present in the base PCP protocol.

### 2.1. PCP PEER Request

A PCP PEER request could contains the destination IP address, port and Transport protocol of the peer the host will be trying to communicate . In that case, if the NAT device maintains a mapping of

zones (and associated NAT pools) to network prefixes it can choose the appropriate NAT pool. It is important to understand that this will only work if the policy that decides to which Zone to route packets is only based on the information present on the PCP PEER request.

Therefore if the PCP Client knows it is behind a NAT with zone support, it is RECOMMENDED that it includes the remote peer's 5-tuple in the PCP PEER request in the connect-then-lifetime case. If the peer's 5-tuple is not present in the PCP request, the external IP and port returned in the message is non-deterministic.

## 2.2. PCP MAP Request

In the case of PCP MAP request the NAT device does not know from which zone to install a mapping and consequently from which NAT pool to choose an external IP address and port. A FILTER Option may be included to allow the PCP Server select the external address pool to use. If other information than the destination IP address is used to drive the selection of the external address pool, additional information is required to be conveyed in the PCP MAP request (e.g., DSCP marking policy (see <http://tools.ietf.org/html/draft-boucadair-pcp-extensions-01#section-3>)).

## 3. PCP Extension for Multiple Zones

The proposed PCP extension is a new PCP Option that would convey the Zone-ID. The Zone-ID is an opaque identifier that is known by the PCP Client and the PCP-controlled NAT device. The procedure to provision the Zone-ID is out of scope.

When the NAT device receives a PCP request with a Zone-ID, it will use that or a derivative of it to determine the NAT pool from which to allocate an IP address and port.

Option Name: ZONEID

Number: TBA (IANA); Mandatory to process

Purpose: It allows the client request and server indicate from which Zone-ID the external IP:port were allocated.

Valid for Opcodes: MAP, PEER

Length: Variable

May appear in: both

Maximum occurrences: 1

#### 4. IANA Considerations

TBD

#### 5. Security Considerations

Subscribers can only request ports for the specific Zone-IDs allowed in their security profile. For example, in a typical Wireless deployment, mobile terminals could request mappings in zones 'Internet', 'HTTP Proxy Farm', and 'Video Farm'. A PCP request that contains a zone-id considered a security violation would be silently dropped.

#### 6. Acknowledgements

Thanks to Mohamed Boucadair for early review comments

#### 7. References

##### 7.1. Normative References

- [RFC0959] Postel, J. and J. Reynolds, "File Transfer Protocol", STD 9, RFC 959, October 1985.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2766] Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.
- [RFC2960] Stewart, R., Xie, Q., Morneault, K., Sharp, C., Schwarzbauer, H., Taylor, T., Rytina, I., Kalla, M., Zhang, L., and V. Paxson, "Stream Control Transmission Protocol", RFC 2960, October 2000.
- [RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.

- [RFC4966] Aoun, C. and E. Davies, "Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status", RFC 4966, July 2007.
- [RFC5382] Guha, S., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", BCP 142, RFC 5382, October 2008.
- [RFC5508] Srisuresh, P., Ford, B., Sivakumar, S., and S. Guha, "NAT Behavioral Requirements for ICMP", BCP 148, RFC 5508, April 2009.

## 7.2. Informative References

- [I-D.ietf-behave-address-format]  
    Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", draft-ietf-behave-address-format-10 (work in progress), August 2010.
- [I-D.ietf-behave-dns64]  
    Bagnulo, M., Sullivan, A., Matthews, P., and I. Beijnum, "DNS64: DNS extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", draft-ietf-behave-dns64-11 (work in progress), October 2010.
- [I-D.ietf-behave-ftp64]  
    Beijnum, I., "An FTP ALG for IPv6-to-IPv4 translation", draft-ietf-behave-ftp64-12 (work in progress), July 2011.
- [I-D.ietf-behave-v6v4-framework]  
    Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", draft-ietf-behave-v6v4-framework-10 (work in progress), August 2010.
- [I-D.ietf-behave-v6v4-xlate-stateful]  
    Bagnulo, M., Matthews, P., and I. Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", draft-ietf-behave-v6v4-xlate-stateful-12 (work in progress), July 2010.
- [I-D.ietf-pcp-base]  
    Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", draft-ietf-pcp-base-16 (work in progress), October 2011.

- [I-D.wing-behave-dns64-config]  
Wing, D., "IPv6-only and Dual Stack Hosts on the Same Network with DNS64", draft-wing-behave-dns64-config-03 (work in progress), February 2011.
- [RFC3261] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [RFC5853] Hautakorpi, J., Camarillo, G., Penfield, R., Hawrylyshen, A., and M. Bhatia, "Requirements from Session Initiation Protocol (SIP) Session Border Control (SBC) Deployments", RFC 5853, April 2010.

#### Author's Address

Reinaldo Penno  
Juniper Networks  
1194 N Mathilda Avenue  
Sunnyvale, California 94089  
USA

Email: rpenno@juniper.net



