

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 27, 2012

M. Boucadair
France Telecom
R. Penno
D. Wing
Cisco
April 25, 2012

Some Extensions to Port Control Protocol (PCP)
draft-boucadair-pcp-extensions-03

Abstract

This document extends Port Control Protocol (PCP) with new functionalities.

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 October 27, 2012.

Copyright Notice

Copyright (c) 2012 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

Internet-Draft

Extensions to PCP

April 2012

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. | DESCRIPTION | 3 |
| 3. | DSCP_POLICY | 4 |
| 4. | CAPABILITY | 5 |
| 5. | REPORT | 8 |
| 6. | CLIENT_IDENTIFIER | 9 |
| 7. | Security Considerations | 10 |
| 8. | IANA Considerations | 11 |
| 9. | Normative References | 11 |
| | Authors' Addresses | 11 |

1. Introduction

This document extends the base PCP [[I-D.ietf-pcp-base](#)] with various PCP Options.

Some of these options may be defined as new PCP OpCodes.

The main goal of this document is to kick-off discussions on the need to define some useful PCP options which are not part of base PCP.

2. DESCRIPTION

This option (Code TBA, Figure 1) MAY be included in a PCP MAP request to include a description associated with a requested mapping. This option is optional to be supported by PCP Servers and PCP Clients. The maximum length SHOULD be a configurable option in the PCP Server. If a PCP Client includes a Description PCP option with a length exceeding the maximum length supported by the PCP Server, only the portion of the Description field fitting that maximum length is stored by the PCP Server.

This option can be used by a user to indicate a description associated with a given mapping such as "My mapping for my FTP server" or "My remote access to my CP router", etc. In addition, in the some deployment scenarios, this field can be used for troubleshooting purposes and can be used to convey values as the ones listed hereafter:

- o "This is the mapping for my specific IPsec implementation"
- o "This is the mapping for subscriber bob@example.com"
- o "This is the mapping for special subscriber
adsl-line-1234@example.com"

- o "This is the mapping that failed before due to XYZ"

Issues related to the usage of this field for troubleshooting or for any further usage are out of scope of this document.

This Option:

Option Name: Description Option (DESCRIPTION)

Number: TBA (IANA)

Purpose: Used to associate a text description with a mapping

Valid for Opcodes: MAP

Length: Variable

May appear in: both request and response

Maximum occurrences: 1




Figure 1: Description Option

3. DSCP_POLICY

In some scenarios, the DSCP marking in the internal interface (i.e., customer-facing interface) and the external one (i.e., Internet-facing interface) of the PCP-controlled device may be distinct. A Service Provider MAY allow its customers to configure their DSCP

marking policies in an upstream device. Distinct DSCP marking policies can be implemented in the internal and external sides of the PCP-controlled device. A PCP Client MAY issue a PCP MAP request indicating its internal DS code point and the external DSCP value. Instructed forwarding policies are applied only for packets marked with a given DSCP value.

A Service Provider may not support DSCP re-marking feature and adopt a transparent scheme to QoS policy enforcement, that is, not controllable by subscribers. Generic QoS enforcement policies can be enforced for all customers: such as leave DSCP field values unchanged.

This option is mandatory to process.

This option (Code TBA, Figure 2) allows to:

- o Re-write any DSCP value to a specific value;

- o Re-write a specific DSCP value to another specific value.

This Option:

Option Name: PCP DSCP Marking Policy Option (DSCP_POLICY)

Number: TBA (IANA); mandatory to process


Purpose: Associated a DSCP re-marking policy with a mapping

Valid for Opcodes: MAP, PEER

Length: 0x04

May appear in: both request and response

Maximum occurrences: 1



DIR : Indicates the direction:

0 Inbound
 1 Outbound
 2 Both
 Int DSCP: Indicates the DSCP value in the customer-faced interface.
 0x3F is used to indicate ANY value.
 Ext DSCP: Indicates the DSCP value in the Internet-faced interface.
 0x3F is used to indicate ANY value.

Figure 2: DSCP Marking option

4. CAPABILITY

The CAPABILITY option (Code: TBA, Figure 3) is used by a PCP Server to indicate to a requesting PCP Client the capabilities it supports with regards to port forwarding operations. Several Capability options MAY be conveyed in the same PCP response message if several functions are co-located in the same PCP-controlled device (e.g., NAT44 and NAT64, NAT44 and ports set assignment capability, etc.).

This option, when received from a PCP Server, is used by a PCP Client to constraint the content of its requests and therefore avoid errors.

This Option:

Option Name: PCP Capabilities Option (CAPABILITY)

Number: TBA (IANA)


Purpose: Retrieve the capabilities of a PCP-controlled device

Valid for Opcodes: can be returned in a error message

Length: 0x01

May appear in: both request and response

Maximum occurrences: None



```

| F T P A S C I O | 00...00 |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 3: Capability option

Below is provided a description of the F, T, P, A, S, C, I and O bits:

| Name | Description |
|------|--|
| F | This bit indicates the address family of the source address issued by internal hosts |
| T | This bit indicates the address family of the source address of the packets forwarded in the external side of the PCP-controlled device |
| P | This bit indicates whether the source port number is translated or not. |
| A | This bit indicates whether the source IP address is translated or not. |
| S | This bit indicates whether the controlled device supports the ability to assign a set of ports |
| C | This bit indicates whether the PCP-controlled devices inspect the received packets and if it can block them |
| I | This bit indicates whether incoming packets are rejected unless an explicit rule is enforced in the PCP-controlled device |
| O | This bit indicates whether outbound packets are inspected or not before being granted to leave the internal realm. |

The value of the F, T, P, A, S, C, I and O bits are as follows:

| Position | Name | Meaning |
|----------|--------------------|--------------------------------|
| 1 | From (F) | 0=from IPv4, 1=from IPv6 |
| 2 | To (T) | 0=to IPv4, 1=to IPv6 |
| 3 | Port-Xlate (P) | 1=translated, 0=not translated |
| 4 | Addr-Xlate (A) | 1=translated, 0=not translated |
| 5 | Port-Set (S) | 1=enabled, 0=not supported |
| 6 | Packet-Control (C) | 1=enabled, 0=not supported |

7 Direction-Out (I) 1=enabled, 0=disabled
8 Direction-In (O) 1=enabled, 0=disabled

A stateless NAT64 [[RFC6145](#)] would have the following values:

From=0 (IPv4)
To=1 (IPv6)
Port-Xlate=0 (No)
Addr-Xlate=1 (Yes)
Port-Set=0 (No)
Packet-control=0 (No)
Direction-out (0) (No)
Direction-In=0 (No)

A stateful NAT64 [[RFC6146](#)] would have the following values:

From=0 (IPv4)
To=1 (IPv6)
Port-Xlate=1 (Yes)
Addr-Xlate=1 (Yes)
Port-Set=0 (No)
Packet-control=0 (No)
Direction-out (0) (No)
Direction-In=0 (No)

A NAT44 would be characterized as follows:


```

To=0 (IPv4)
Port-Xlate=1 (Yes)
Addr-Xlate=1 (Yes)
Port-Set=0 (No)
Packet-control=0 (No)
Direction-out (0) (No)
Direction-In=0 (No)

```

5. REPORT

The Report PCP Option (Code TBA, Figure 4) is used by a PCP Client to report a set of useful information to the PCP Server. Several Report Options with distinct Report Sub-Code values MAY be conveyed in the same PCP message. Only report data associated with the PCP Server to which this option is sent MUST be included in a Report Option.

This option can be used for troubleshooting or diagnose purposes.

This Option:

```

Option Name: PCP Report Option (REPORT)
Number: TBA (IANA)
Purpose: Send a set of report data
Valid for Opcodes: MAP
Length: Variable
May appear in: both request and response
Maximum occurrences: Multiple

```




Figure 4: Report Option

The following Report Sub-Code values are defined:

| Position | Meaning |
|----------|---------|
|----------|---------|

| | |
|------|---|
| 0x00 | Time since last reboot/boot |
| 0x01 | Count of transmitted PCP messages to the PCP Server since last boot |
| 0x02 | Count of retransmitted PCP messages to the PCP Server since last boot |
| 0x03 | Count of received PCP Error messages from the PCP Server |

6. CLIENT_IDENTIFIER

PCP CLIENT_ID (Code TBA, Figure 5) is a token randomly [[RFC4086](#)] generated by the PCP Client. Only one CLIENT_ID Option MUST be present in a PCP message. The PCP Client and PCP Server MUST store the value included in this Option in a PCP MAP request.

- o The CLIENT_ID MUST be generated by the PCP Client and not the PCP Server;
- o Upon change of the IP address of the PCP Client (or a third party on behalf of which a mapping has been created), the CLIENT_ID is used to update related mappings (e.g., PCP MAP delete request and PCP MAP create request);
- o The same CLIENT_ID MUST be used for all requested mappings, unless a new CLIENT_ID is generated by the PCP Client (e.g., reboot, OS crash, etc.);
- o The CLIENT_ID is stored by the PCP Server for all mappings (persistent storage);

This Option:

Option Name: PCP Client Identifier Option (CLIENT_ID)

Number: TBA (IANA); mandatory to process option

Purpose: Associate an identifier with the mappings

Valid for Opcodes: MAP

Length: Variable

May appear in: both request and response

Maximum occurrences: 1

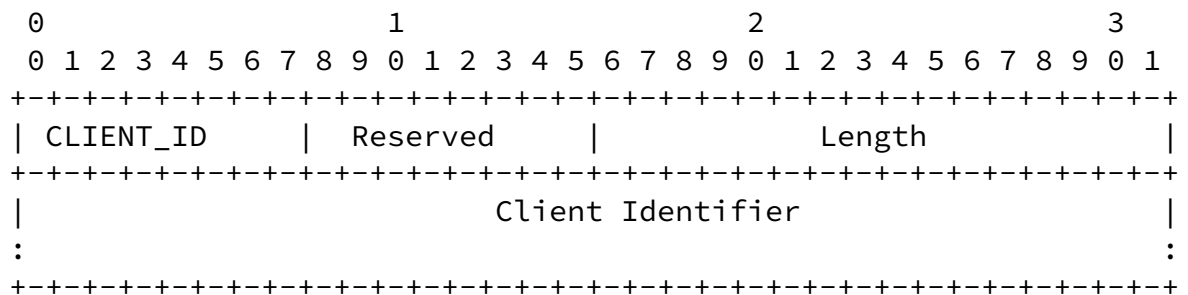


Figure 5: CLIENT_ID PCP Option

The length of the CLIENT_ID is encoded in the Length field in bytes. The length of the CLIENT_ID MUST be at least 4 bytes and MUST NOT exceed 16 bytes.

The RECOMMENDED value is 16 bytes so as to have a robust random CLIENT_ID. If a CLIENT_ID longer than 16 bytes or shorter than 4 bytes is received, the PCP Server MUST issue a PCP Error message with an error cause equal to "Invalid Client-ID".

For sanity checks, a PCP Server maintains the same CLIENT_ID value (which is used in the latest PCP request) for a given PCP Client for all mappings associated with the same internal IP address belonging to the same subscriber. Indeed, the PCP Server maintains an additional identifier denoted as subscriber-Id. A subscriber-Id can be an IP address, IPv6 prefix or a subscriber identifier configured locally.

[7.](#) Security Considerations

Security considerations discussed in [[I-D.ietf-pcp-base](#)] must be considered. The use of CLIENT_ID option allows to soften issues related to stale mappings.

Boucadair, et al.

Expires October 27, 2012

[Page 10]

Internet-Draft

Extensions to PCP

April 2012

[8.](#) IANA Considerations

The following PCP Option Codes are to be allocated:

DESCRIPTION

DSCP_POLICY: The "0" bit MUST be set to 1.

CAPABILITY

REPORT

CLIENT_IDENTIFIER: The "0" bit MUST be set to 1.

[9.](#) 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-24](#) (work in progress), March 2012.

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

[RFC4086] Eastlake, D., Schiller, J., and S. Crocker, "Randomness Requirements for Security", [BCP 106](#), [RFC 4086](#), June 2005.

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", [RFC 6145](#), April 2011.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful

NAT64: Network Address and Protocol Translation from IPv6
Clients to IPv4 Servers", [RFC 6146](#), April 2011.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes, 35000
France

Email: mohamed.boucadair@orange.com

Boucadair, et al. Expires October 27, 2012 [Page 11]

Internet-Draft Extensions to PCP April 2012

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

