

PCP Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 23, 2015

M. Boucadair
France Telecom
October 20, 2014

PCP as a Traffic Classifier Control Protocol
draft-boucadair-pcp-sfc-classifier-control-01

Abstract

This document specifies how PCP (Port Control Protocol) can be used as a classifier control protocol.

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 23, 2015.

Copyright Notice

Copyright (c) 2014 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	2
2.	Scope of this Document	3
3.	Objectives for Controlling Classifiers	3
4.	PCP as a Traffic Classifier Control Protocol	4
5.	Deployment Model	5
6.	Missing PCP Extensions	5
7.	Detailed Specification	6
8.	IANA Considerations	6
9.	Security Considerations	6
10.	Acknowledgements	6
11.	References	6
11.1.	Normative References	6
11.2.	Informative References	6
	Author's Address	7

[1.](#) Introduction

PCP (Port Control Protocol, [[RFC6887](#)]) has been specified to control an upstream function such as NATs or firewalls. PCP can be used to interact with both statefull and stateless functions.

PCP can be abstracted as a means to notify an upstream network with the flow characteristics that would trigger decisions on the appropriate policies to be applied on these flows at the network side. This document focuses on a typical function that is present in operational networks: Traffic Classifier Function. Two examples are listed below:

DiffServ Classifiers A typical example of packet classifier is DiffServ Classifier [[RFC2475](#)] that is responsible to select packets in a traffic stream based on the content of some portion of the packet header: this can be based solely on the DSCP field or based on a combination of more header fields, such as source address, destination address, DSCP field, protocol ID, source port and destination port numbers, and other information such as incoming interface. These classifiers must be configured by some means as documented in [[RFC2475](#)]:

"Classifiers are used to "steer" packets matching some specified rule to an element of a traffic conditioner for further processing. Classifiers must be configured by some

management procedure in accordance with the appropriate TCA (Traffic Conditioning Agreement)."

This document proposes PCP to configure DiffServ Classifiers.

SFC Classifiers Another classifier is SFC Classifier (e.g., [[I-D.ietf-sfc-problem-statement](#)] [[I-D.ietf-sfc-architecture](#)]). This classifier is responsible for classifying flows to determine which Service Function Chain (SFC) they belong to.

Similar to DiffServ, an SFC Classifier can rely on a variety of classifying rules.

PCP can be used to instruct SFC Classifier policies.

A traffic classifier (or classifier for short) is a function that is responsible for classifying flows based on (pre-defined) rules. Once the traffic is classified, it can be marked to bear its class of service (DSCP re-marking [[RFC2474](#)]), dropped, shaped, or any other action instructed by the matching rule. This document focuses on classification rules that manipulate L3/L4 fields of IP packets.

A Classifier can be seen as a statefull service function that applies a set of policies for packets and/or flows matching a set of criteria. This document specifies how PCP can be used as a classifier control protocol.

Note a classifier can be co-located with a CGN (Carrier Grade NAT, [[RFC6888](#)]), or a firewall. PCP can be used to install policies in all these functions.

[2.](#) Scope of this Document

This version of the document explains the motivations, basic assumptions, and identifies some missing features. Detailed specification of required extensions will be elaborated in future

versions of the document.

This document focuses on the control of L2/L3/L4 Classifiers. Sophisticated classifiers based on heuristics (e.g., those involving DPI (Deep Packet Inspection) modules) are out of scope.

[3.](#) Objectives for Controlling Classifiers

Below are listed some objectives for a classifier control means:

- o Rationalize the management of classification rules.

Boucadair

Expires April 23, 2015

[Page 3]

Internet-Draft

PCP & Traffic Classification

October 2014

- o Help assessing the impact of removing or modifying a classification rule.
- o Check the coherency of instantiated classification rules.
- o Help aggregating rules: this allows to optimize the classification rule table and therefor accelerate packet/flow processing (mainly reduce lookup delays).
- o Adjust classification rules when rules are based on volatile identifiers (e.g., IP address).
- o Maintain an global overview of instantiated rules in involved Network Elements.
- o Rapidly restore state during failure events.
- o Network Elements can retrieve their table after failure events without requiring permanent storage capacity.

[4.](#) PCP as a Traffic Classifier Control Protocol

PCP fulfils most of the objectives listed in the previous section. Concretely, PCP allows for the following:

- o Directionality
- o Classification involving port sets

- o Classification rules involving IPv4 addresses
- o Classification rules involving IPv4 prefixes
- o Classification rules involving IPv6 addresses
- o Classification rules involving IPv6 prefixes
- o Classification rules with or without NAT
- o Feedback loop to assess whether a classification rule has been successfully enforced: PREFER_FAILURE
- o Associate a description with classification rules
- o Classification rules are associated with lifetimes
- o PCP client can re-install states if a loss is detected

- o PCP server does not need to store the entries in a permanent storage; state can be installed by the PCP client
- o PCP client detects rapidly any state loss, including reboot of the PCP server
- o Multiple PCP clients can control the same PCP server
- o 2-way exchange is needed to install new state
- o No need to lock the server waiting when concurrent clients are in contact with the server.

[5.](#) Deployment Model

The reference architecture adopted in this document assumes that both the PCP client and server are managed by the same administrative entity (e.g., an operator).

Classification rules are not exposed outside an administrative domain. In particular, subscribers are not aware of these policies.

PCP requests received in the subscriber-faced interfaces are not allowed to manage policies enforced in the classifiers.

6. Missing PCP Extensions

Some candidate extensions are listed below:

- o Extended THIRD_PARTY option to include a L2 identifier (e.g., MAC address), an opaque subscriber-identifier, an IMSI, etc. A candidate option is defined in [[I-D.ripke-pcp-tunnel-id-option](#)].

This extension is also needed to control NATs that are Layer-2 Aware (see [Section 2.1 of \[RFC6887\]](#)).

- o Extended FILTER to include a remote range of ports. This extension might be useful for the firewall case.
- o DSCP-based filtering. This extension is also useful for the firewall control case.
- o DSCP re-marking: this is to be enforced at the boundaries of a domain. The marking at the access segment may not be the same as the core network. Marking must be enforced by a trusted device.
- o Explicit actions: re-mark, drop, shape (can be used with FLOWDATA [[I-D.wing-pcp-flowdata](#)]).

Boucadair

Expires April 23, 2015

[Page 5]

Internet-Draft

PCP & Traffic Classification

October 2014

- o A means to indicate the SFC binding.

7. Detailed Specification

This section will be completed if the working group agrees with the problem to be solved.

8. IANA Considerations

To be completed.

9. Security Considerations

Security considerations discussed in [[RFC6887](#)] MUST be taken into account.

[10.](#) Acknowledgements

TBC.

[11.](#) References

[11.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC6887] Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", [RFC 6887](#), April 2013.

[11.2.](#) Informative References

- [I-D.ietf-sfc-architecture]
Halpern, J. and C. Pignataro, "Service Function Chaining (SFC) Architecture", [draft-ietf-sfc-architecture-02](#) (work in progress), September 2014.
- [I-D.ietf-sfc-problem-statement]
Quinn, P. and T. Nadeau, "Service Function Chaining Problem Statement", [draft-ietf-sfc-problem-statement-10](#) (work in progress), August 2014.
- [I-D.ripke-pcp-tunnel-id-option]
Ripke, A., Dietz, T., Quittek, J., and R. Silva, "PCP Tunnel-ID Option", [draft-ripke-pcp-tunnel-id-option-01](#) (work in progress), July 2014.

- [I-D.wing-pcp-flowdata]
Wing, D., Penno, R., and T. Reddy, "PCP Flowdata Option", [draft-wing-pcp-flowdata-00](#) (work in progress), July 2013.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", [RFC 2474](#), December 1998.

- [RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", [RFC 2475](#), December 1998.
- [RFC6888] Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", [BCP 127](#), [RFC 6888](#), April 2013.

Author's Address

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com