

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 4, 2011

M. Boucadair
France Telecom
S. Sivakumar
Cisco
March 03, 2011

Reservation of Pair of Ports with PCP
draft-boucadair-pcp-rtp-rtcp-01

Abstract

This document defines a new PCP Option to reserve a pair of ports in a PCP-controlled device while preserving the parity and contiguity. This new PCP Option eases the traversal of NAT for RTP/RTCP flows.

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 September 4, 2011.

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
2. Why this PCP Option is Needed?	3
3. Definition of the Port Reservation Option	4
3.1. Requirements	4
3.2. Rationale	4
3.3. PCP Port Reservation Option	5
4. Client Behaviour	5
5. Server Behaviour	6
6. Illustration Examples	7
6.1. Port Reservation Option Not Supported by The PCP Server .	7
6.2. Port Reservation Option Is Supported by The PCP Server .	8
6.3. Delete the Mappings	10
7. IANA Considerations	12
8. Security Considerations	12
9. References	12
9.1. Normative References	12
9.2. Informative References	13
Authors' Addresses	13

1. Introduction

This document defines a new PCP Option [I-D.ietf-pcp-base] which aims to ease the traversal of RTP/RTCP based applications [RFC3550] when a NAT is involved in the path.

The main advantage of using PCP is it does not need any further feature to be supported by the outbound proxy to assist the remote endpoint to successfully establish media sessions. In particular, there is no need to implement what is commonly denoted as HNT (hosted NAT traversal) or any keepalive scheme for media streams. Furthermore, ALGs are not required in the NAT for this purpose.

Note that the base PCP allows to retrieve the external IP address and port to be conveyed in the SIP signaling messages [RFC3261]. Therefore SIP Proxy Servers can be more lightweight and do not need to support means to ease the NAT traversal of SIP messages (e.g., [RFC5626], [I-D.ietf-sipcore-keep], etc.). The advantage of using the external IP address and port is this provides a hint to the proxy server there is no need to return a small expire timer (e.g., 60s).

It is worth mentioning that based on some deployments, the impact of HNT in the proxy servers is very severe: when this function is enabled, a device dimensioned to service 75k users can serve only 30k users!

2. Why this PCP Option is Needed?

Traditionally the voice/video applications that use RTP and RTCP would specify only the RTP port that the application would use for streaming the RTP data. The inherent assumption is that the RTCP traffic will be sent on the next higher port. Below is provided an excerpt from [RFC3550]:

" RTP relies on the underlying protocol(s) to provide de-multiplexing of RTP data and RTCP control streams. For UDP and similar protocols, RTP SHOULD use an even destination port number and the corresponding RTCP stream SHOULD use the next higher (odd) destination port number. For applications that take a single port number as a parameter and derive the RTP and RTCP port pair from that number, if an odd number is supplied then the application SHOULD replace that number with the next lower (even) number to use as the base of the port pair. For applications in which the RTP and RTCP destination port numbers are specified via explicit, separate parameters (using a signaling protocol or other means), the application MAY disregard the restrictions that the port numbers be even/odd and consecutive although the use of an even/

odd port pair is still encouraged."

[RFC3605] defines an explicit "a=RTCP" SDP attribute for some applications using a distinct port than RTP+1. Even though [RFC3605] defines a new attribute for explicitly specifying the RTCP attribute for the SDP based applications, but since it is not a MUST to use this attribute, there are still applications that are not compliant with this RFC. There are also non-SDP based applications that use RTP/RTCP like H323, that make the assumption that RTCP streaming will happen on RTP+1 port.

In order for these applications to work across NAT, the NAT device must have an application layer gateway, that would allocate two consecutive ports. In a PCP context, a similar functionality need to be provided for the PCP Client to request two consecutive ports and the PCP Server to allocate and respond with the information of the allocated port.

This document describes the mechanism to request a pair of consecutive ports for a PCP-controlled device and the corresponding mechanism for the PCP Server to allocate and respond to the port allocation request.

3. Definition of the Port Reservation Option

3.1. Requirements

The PCP Option used to reserve a port pair should meet the following requirements:


1. Preserve the port parity as discussed in Section 4.2.2 of [RFC4787].
2. Preserve port contiguity as discussed in Section 4.2.3 of [RFC4787] (i.e., RTCP=RTP+1).

3.2. Rationale

Since PCP does not support a mechanism to include multiple port numbers in the same request/response, only the RTP port is explicitly signaled in PCP messages. The companion port (i.e., RTCP port) is reserved too when the PCP Server returns back the Port Reservation Option in the response.

3.3. PCP Port Reservation Option

The format of the PCP Port Reservation Option is defined in Figure 1.



This Option:

```
name: Port Reservation Option (PORT_RESRV_OPT)
```

```
number: TBA (IANA)
```

purpose: Used to retrieve a pair of ports

is valid for OpCodes: MAP4, MAP6

```
length: 0
```

may appear in: both

```
maximum occurrences: 1
```

Figure 1: Port Reservation Option

4. Client Behaviour

To retrieve a pair of ports following the requirements listed in Section 3.1, the PCP Client adds the Port Reservation Option to its PCP MAPy (y=4 or 6) request. The PCP Client MAY indicate its preferred RTP external port. This port is likely to be equal to the internal port indicated in the PCP request. A PCP Client is encouraged to use randomized external ports to receive/send its media flows.

Once a response is received from the PCP Server, the PCP Client checks whether the Port Reservation Option is supported by the peer PCP Server following the procedure defined in Section 5.3 of [I-D.ietf-pcp-base].

- o If the answer is positive, the PCP Client retrieves the mapping returned by the PCP Server; in particular the external port number

should be even. This port is indicated to the remote peer as the port number used for RTP flows. RTCP is assumed to use the returned external port number + 1.

- o If the Option is not supported by the PCP Server, and according the port quota, only the RTP port can be signaled to the remote endpoint (e.g., SDP offer/answer [RFC4566]). RTCP flows are likely to fail if no mechanism to assist the traversal of RTCP flows is supported (e.g., "a=RTCP" attribute).

When a pair of ports is retrieved from the PCP Server, two mappings are instantiated in both the PCP Server and PCP Client. For explicit deletion of these mappings, the PCP Client and PCP Server follow the procedure defined in Section 8 of [I-D.ietf-pcp-base] for each mapping.

To reduce the delay to establish media sessions, the PCP Client MAY reserve a pair of ports once the (SIP) registration phase has been successfully completed. These pair of ports will be included in SDP offers/answers for instance.

5. Server Behaviour

Upon receiving the Port Reservation Option, the PCP Server validates the request for the supported OpCode values. If an unrecognized value is received a Invalid request error is returned to the PCP Client (e.g., using MALFORMED_REQUEST error). The reason for rejecting the request could be an invalid internal IP address, invalid Internal port, etc.

For a valid request, the PCP Server collects the Internal port and the hinted external port and verify against any administrative rules to allow or disallow the PCP Client from making this request. An example of an administrative rule will be by fulfilling the request it would put the client over its administratively allowed limits. In those cases, the PCP Server will treat this as an error and this is handled the same way as described in [I-D.ietf-pcp-base] for the denial of honoring the request with the appropriate Opcode.

To handle the PCP Reservation Option by the PCP Server, the procedure defined in Section 5.3 of [I-D.ietf-pcp-base] should be followed. When PCP Reservation Option is not supported, the PCP Server MUST treat the request as any PCP request to create an individual mapping. If port parity preservation is supported by the PCP Server, an even port is likely to be returned to the PCP Client. Otherwise, a port is returned if the port quota is not reached.

The following describes the behavior of the PCP Server when the PCP Reservation Option is supported.

The PCP server should request the controlling NAT device to allocate a pair of consecutive ports. If there is a hinted external port present in the request, the server MAY try to honor the request. The PCP Server MUST honor the parity by requesting the allocation of ports that match the parity. However, there is no guarantee that the hinted external ports are available or be allocated. Two mappings are therefore instantiated by the PCP Server with the same lifetime value. These mappings are treated as any individual mapping.

If the port allocation failed either because of the unavailability of ports or the port parity could not be honored, the PCP Server SHOULD reserve only one external port. The PCP Server SHOULD indicate in the response that the PCP Reservation Option has not been honored as specified in Section 5.3 of [I-D.ietf-pcp-base].

6. Illustration Examples

This section provides a list of examples to illustrate the usage of PCP Port Reservation Option.

6.1. Port Reservation Option Not Supported by The PCP Server

Figure 2 shows an example of the flow exchange which is observed when the PORT_RESERVATION_OPTION is not supported by the PCP Server.




Figure 2: Flow Example of a PCP Server which does not support the Port Reservation Option

6.2. Port Reservation Option Is Supported by The PCP Server

Figure 3 and Figure 4 illustrate two examples of the flow exchanges which are observed when the `PORT_RESERVATION_OPTION` is supported by the PCP Server. Figure 3 shows an example of a PCP Server supporting the option and honoring the requested external port number. Figure 4 shows an example of a PCP Server supporting the option but not honoring the requested external port number.

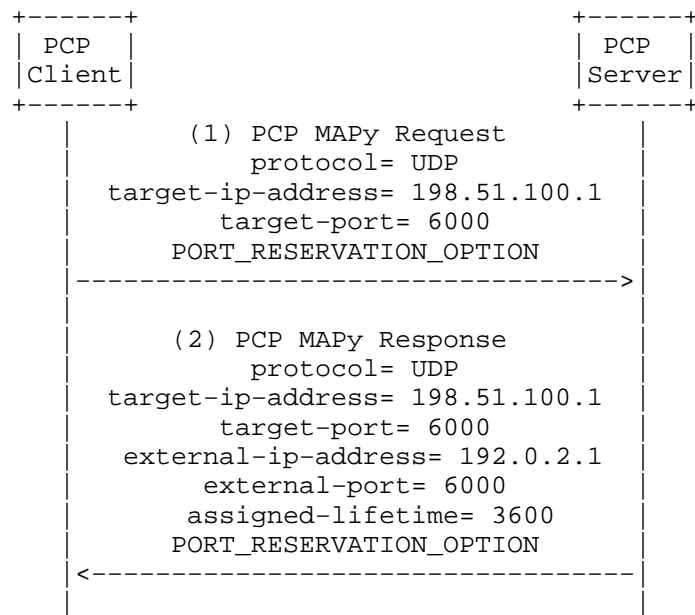


Figure 3: Flow Example of a PCP Server supporting the option and honoring the hinted external port

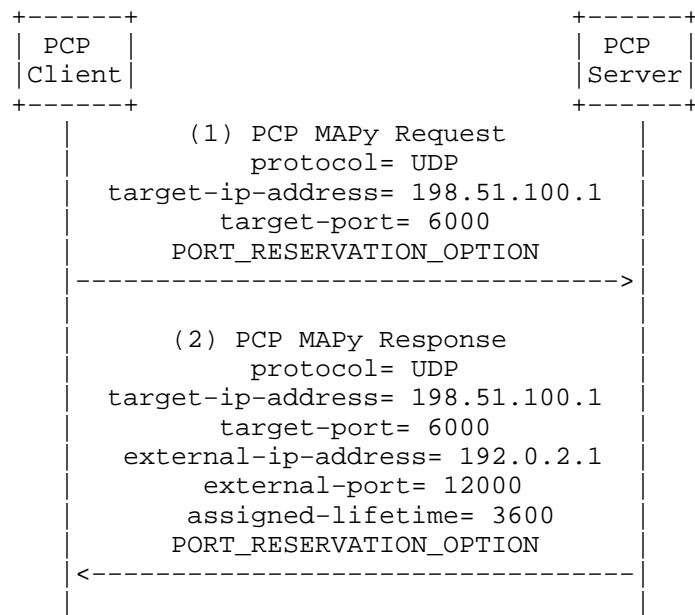


Figure 4: Flow Example of a PCP Server supporting the option but not honoring the hinted external port

6.3. Delete the Mappings

Figure 5 and Figure 6 shows the exchanges that occur to delete the created mappings.

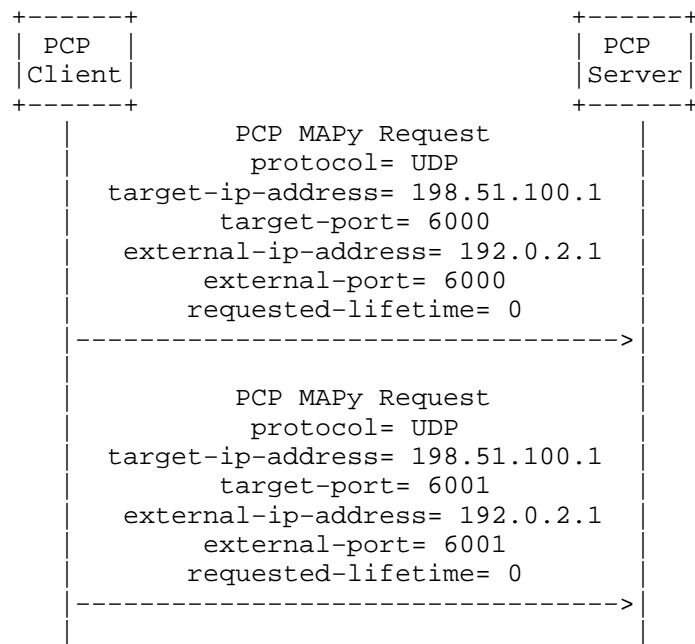


Figure 5: Flow example to delete the mappings

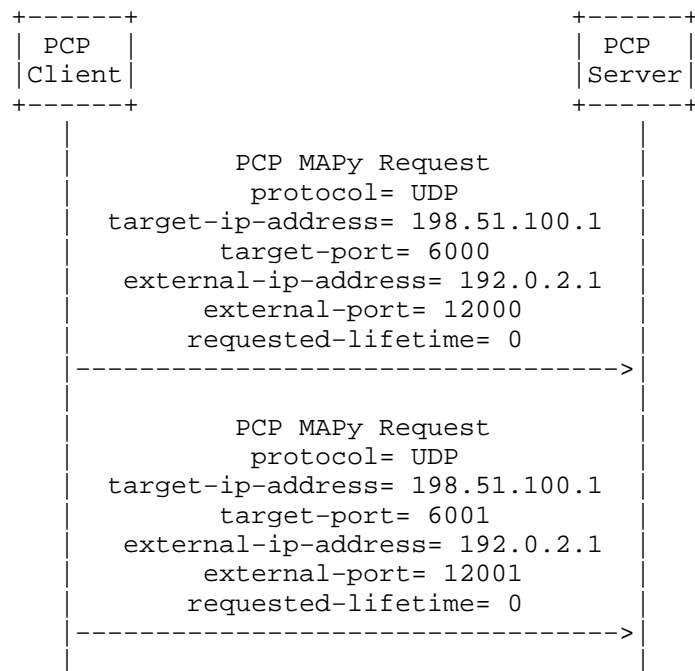


Figure 6: Flow example to delete the mappings (2)

7. IANA Considerations

This document request the assignment of a new PCP Option code:

- o PORT_RESERVATION_OPTION.

8. Security Considerations

This document does not introduce any security issue in addition to what is taken into account in [I-D.ietf-pcp-base].

9. References

9.1. Normative References

[I-D.ietf-pcp-base]
 Wing, D., Cheshire, S., Boucadair, M., Penno, R., and F. Dupont, "Port Control Protocol (PCP)",
 draft-ietf-pcp-base-06 (work in progress), February 2011.

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [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.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [RFC3605] Huitema, C., "Real Time Control Protocol (RTCP) attribute in Session Description Protocol (SDP)", RFC 3605, October 2003.
- [RFC4566] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [RFC4787] Audet, F. and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, January 2007.

9.2. Informative References

- [I-D.ietf-sipcore-keep] Holmberg, C., "Indication of support for keep-alive", draft-ietf-sipcore-keep-12 (work in progress), January 2011.
- [RFC5626] Jennings, C., Mahy, R., and F. Audet, "Managing Client-Initiated Connections in the Session Initiation Protocol (SIP)", RFC 5626, October 2009.

Authors' Addresses

Mohamed Boucadair
France Telecom
Rennes, 35000
France

Email: mohamed.boucadair@orange-ftgroup.com

Senthil Sivakumar
Cisco
7100 Kit Creek Road
Research Triangle Park, North Carolina 27709
USA

Email: ssenthil@cisco.com

