

PCP NAT64 Experiments

I-D. draft-boucadair-pcp-nat64-experiments

IETF 85-Atlanta, November 2012

Authors: M. Ait Abdesselam,
M. Boucadair, A. Hasnaoui, J. Queiroz

Presenter: J. Queiroz

Objectives of this effort

- Test PCP in IPv6-only environments including WiFi and 3G by:
 - Implement a PCP Server controlling a NAT64 device
 - Implement a PCP Client to be embedded in IPv6-only hosts
 - Update some applications to be PCP-enabled
- Validate the overall PCP chain and assess its impact

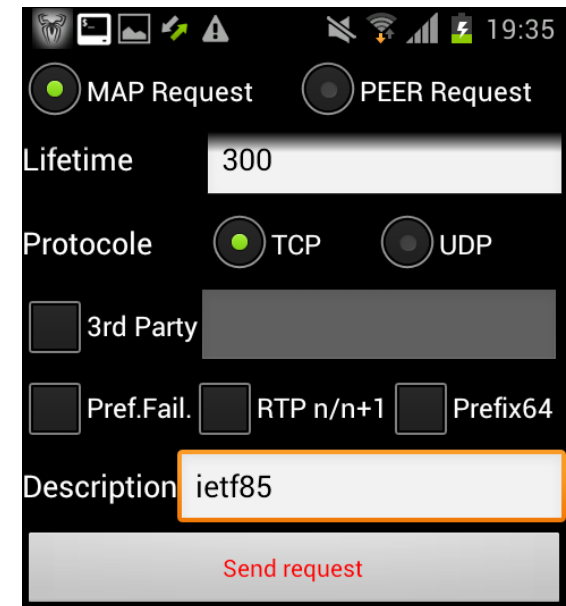

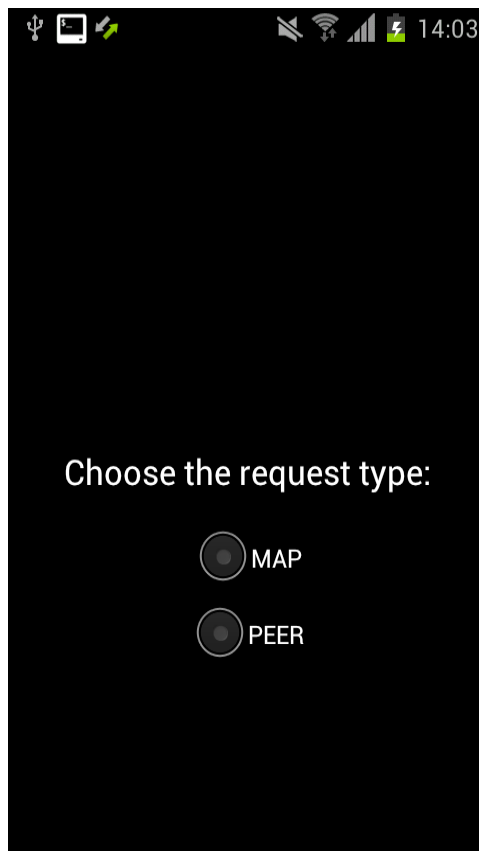
Software Modules

- NAT64
 - Based on Viagenie's Ecdysis implementation (Linux Kernel module).
 - Added a management interface & PCP network interface to be connected via TCP.
- PCP server
 - Based on ISC's PCP server for DS-Lite.
 - Added an interactive shell to view/add/delete users & mappings.
 - Added Description & RTP/RTCP options.
- PCP packet generator for Android
 - Custom packet creator for MAP & PEER Opcodes.
- Linphone SIP UA integrating PCP Client
 - Unique Android SIP client supporting IPv6.
 - Ported from Linux C coded implementation.

Implemented Drafts

- draft-ietf-pcp-base
 - MAP and PEER OpCodes
 - THIRD_PARTY option
 - PREFER_FAILURE option
- I-D.boucadair-pcp-description-option
- I-D.boucadair-pcp-nat64-prefix64-option
- I-D.boucadair-pcp-rtp-rtcp
- I-D.boucadair-mmusic-altc

PCP Packet Generator GUI



```
PCP Server: 2001:688:1f94:3000::2 5351
Internal: 2001:df8:0:112:8d6d:fcf7:e202:108e.8080
Sugg. external addr: 161.105.194.14:8080
Req. Lifetime =300 Seconds | Proto: TCP
Sending ...Request sent successfully!!
Waiting for response ...Received a :60bytes response.
=====> Mapping created <=====
Lifetime= 300
Epoch time =505
assigned external= 161.105.194.14:8080
Description= ietf85
Prefix64=/0
```

PCP commands

```
PCPSERVER> list user
```

```
<===== USERS LISTING =====>
Subscriber(client) ip:<2001:df8:0:112:8d6d:fcf7:e202:108e>
      External IPv4 = <161.105.194.14>
      Port range:[ 1024 - 64999 ]
      Epoch time value =      <2099>
<===== 0 =====>
```

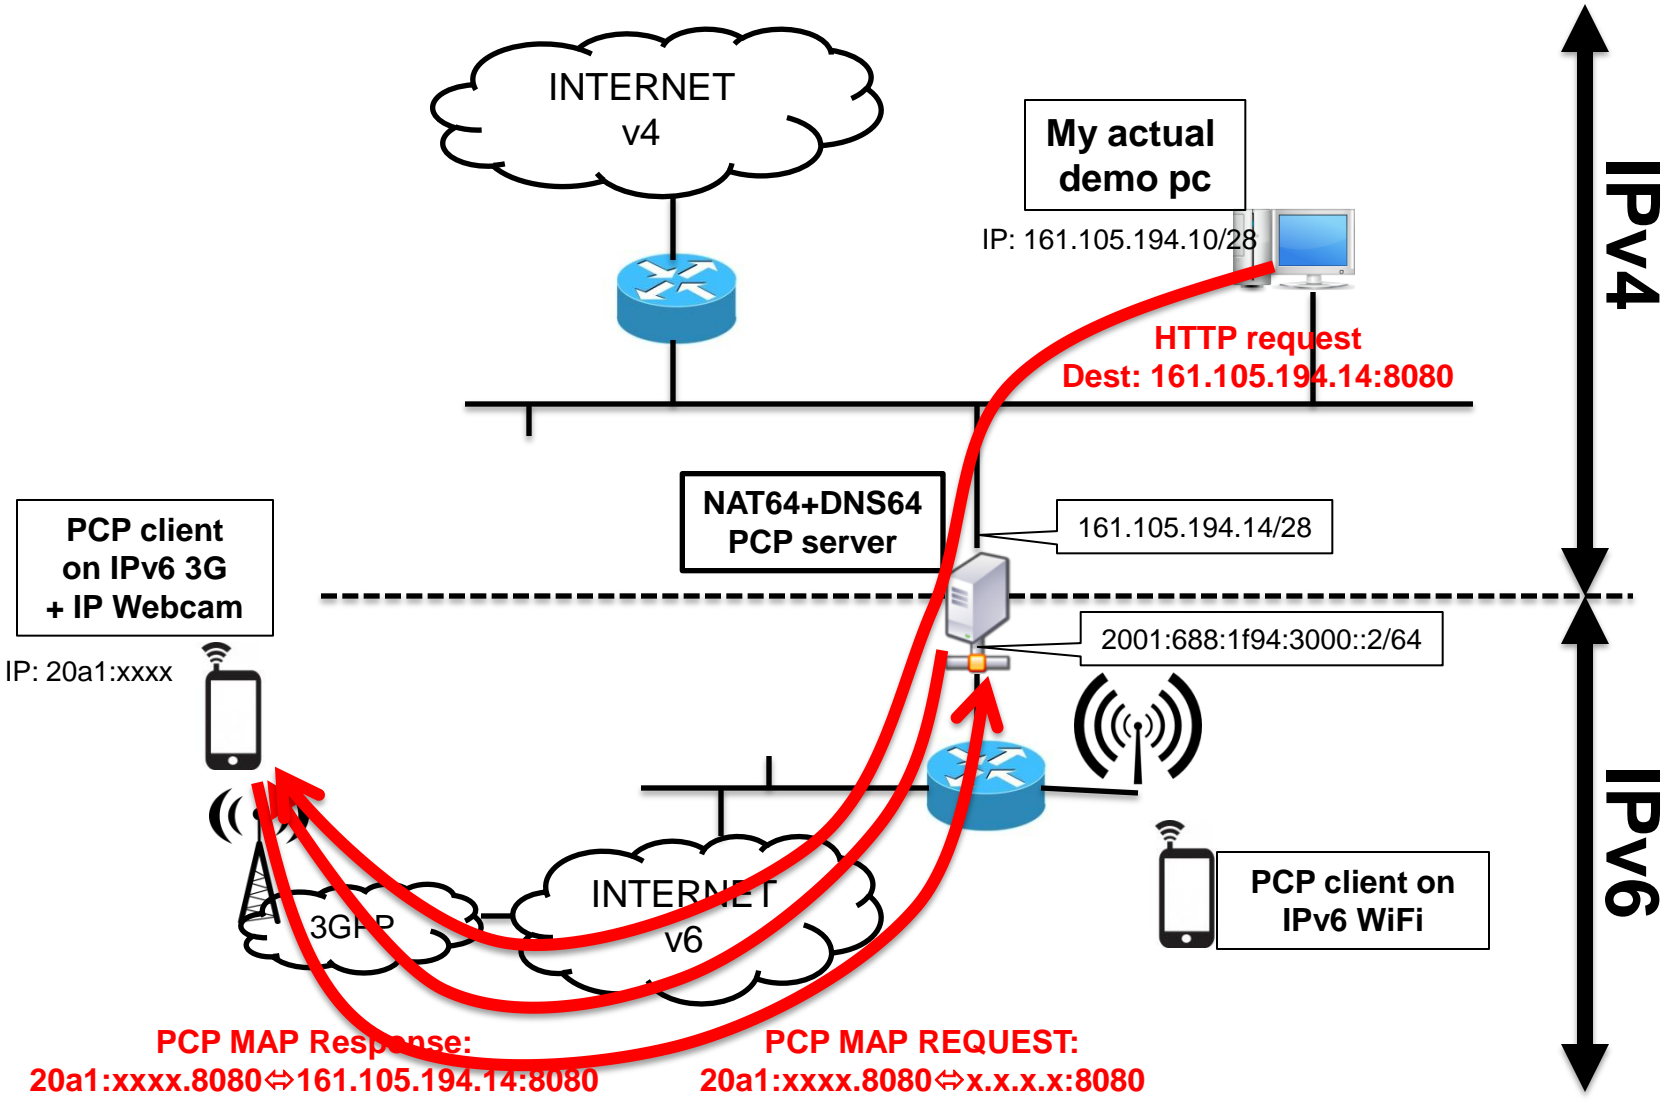
```
PCPSERVER> list map
```

```
<===== MAPPING LISTING BY USER =====>
Subscriber(client) ip: <2001:df8:0:112:8d6d:fcf7:e202:108e>
      External IPv4 = <161.105.194.14>
      Port range:      [ 1024 - 64999 ]
      User epoch time=      <2102>
      ---- List of mappings : ----
          protocole=      TCP
          Internal port=   <8080>
          External port = <8080>
          Expire= <2474>
          Description=    "ietf85"
          No 3rd party address

          -----○-----
<===== 0 =====>
```

Test #1: Hosting a Server behind a NAT64

Test #1: Testbed architecture



PCP request

The screenshot displays a Linux desktop environment with a terminal window titled "NX - poiral@161.105.194.14:1027 - NAT64" and a Wireshark window titled "eth3 [Wireshark 1.8.0 (SVN Rev Unknown from unknown)]".

The terminal window shows a Wireshark capture of network traffic. The packet list pane contains the following entries:

No.	Time	Source	Destination	Protocol	Length	Info
5	3.075752000	2001:df8:0:112:8d6d:fcf7:e202:108e	2001:688:1f94:3000::2	PCP	122	Source port: 42495 Destination port: nat-pmp
6	3.076796000	2001:688:1f94:3000::2	2001:df8:0:112:8d6d:fcf7:e202:108e	PCP	122	Source port: nat-pmp Destination port: 42495
7	3.795016000	fe80::20f:3dff:fef4:4319	ff02::1	ICMPv6	134	Router Advertisement from 00:0f:3d:f4:43:19

The packet details pane for the selected packet (No. 6) shows the following structure:

- Frame 5: 122 bytes on wire (976 bits), 122 bytes captured (976 bits) on interface 0
- Ethernet II, Src: Cisco_d0:1b:c9 (00:1b:90:d0:1b:c9), Dst: D-Link_f4:43:19 (00:0f:3d:f4:43:19)
- Internet Protocol Version 6, Src: 2001:df8:0:112:8d6d:fcf7:e202:108e (2001:df8:0:112:8d6d:fcf7:e202:108e), Dst: 2001:688:1f94:3000::2 (2001:688:1f94:3000::2)
- User Datagram Protocol, Src Port: 42495 (42495), Dst Port: nat-pmp (5351)
- Port Control Protocol
 - Version: 1
 - R bit: Request (0)
 - Opcode: MAP (0x01)
 - Reserved: 0
 - Requested Lifetime: 3000 sec
 - PCP Client's IP Address: 2001:df8:0:112:8d6d:fcf7:e202:108e (2001:df8:0:112:8d6d:fcf7:e202:108e)
 - MAP Request
 - Protocol: TCP (6)
 - Reserved: 0
 - Internal Port: 8080
 - Suggested External Port: 0
 - Suggested External IP Address: ::ffff:0.0.0.0 (::ffff:0.0.0.0)
 - Description Option
 - Option Code: description (0x40)
 - Reserved: 0
 - Option Length: 8 bytes
 - Data: 6965746638350000

The packet bytes pane shows the raw data of the PCP request, including the Ethernet II header, IPv6 header, and the PCP payload. The data field of the Description Option is highlighted in the terminal window.

PCP response

The screenshot shows a Linux desktop environment with a terminal window titled "NX - poiral@161.105.194.14:1027 - NAT64" and a Wireshark window titled "eth3 [Wireshark 1.8.0 (SVN Rev Unknown from unknown)]".

The Wireshark packet list shows the following entries:

No.	Time	Source	Destination	Protocol	Length	Info
5	3.075752000	2001:df8:0:112:8d6d:fcf7:e202:108e	2001:688:1f94:3000::2	PCP	122	Source port: 42495 Destination port: nat-pmp
6	3.076796000	2001:688:1f94:3000::2	2001:df8:0:112:8d6d:fcf7:e202:108e	PCP	122	Source port: nat-pmp Destination port: 42495
7	3.795016000	fe80::20f:3dff:fef4:4319	ff02::1	ICMPv6	134	Router Advertisement from 00:0f:3d:f4:43:19

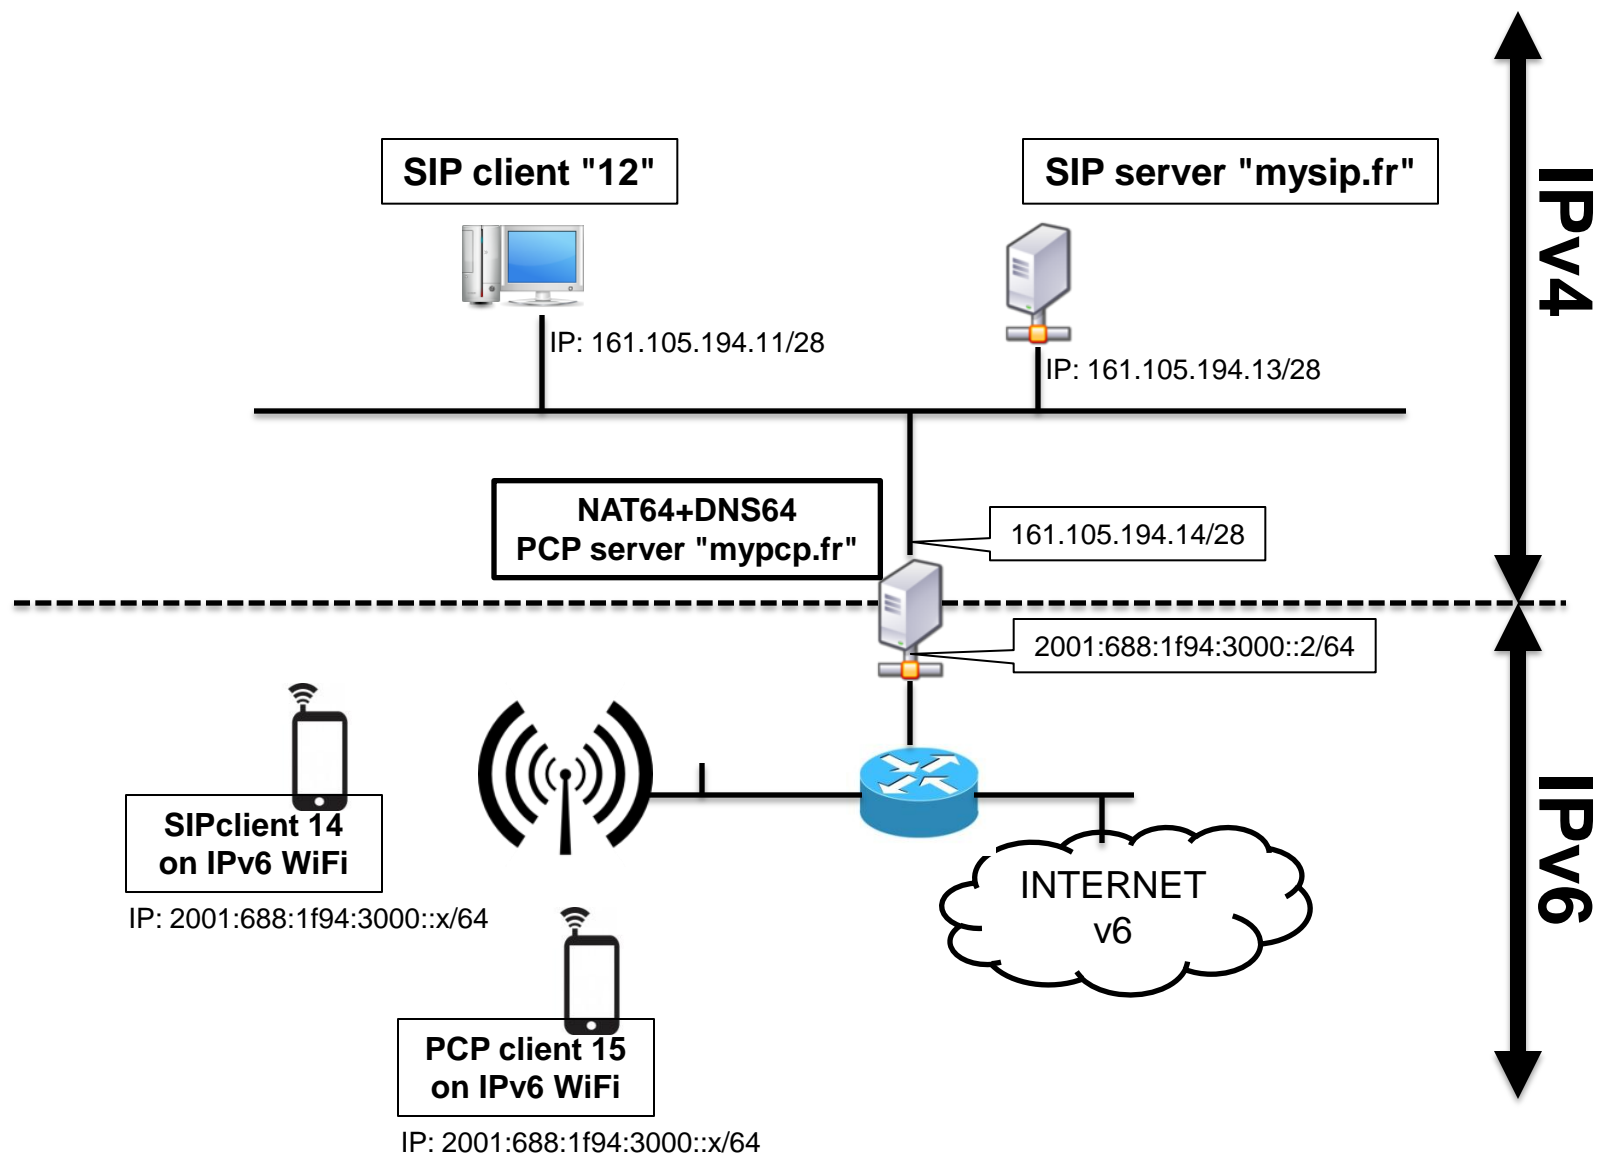
The packet details pane for the selected PCP packet (No. 6) shows the following structure:

- Ethernet II, Src: D-Link_f4:43:19 (00:0f:3d:f4:43:19), Dst: Cisco_d0:1b:c9 (00:1b:90:d0:1b:c9)
- Internet Protocol Version 6, Src: 2001:688:1f94:3000::2 (2001:688:1f94:3000::2), Dst: 2001:df8:0:112:8d6d:fcf7:e202:108e (2001:df8:0:112:8d6d:fcf7:e202:108e)
- User Datagram Protocol, Src Port: nat-pmp (5351), Dst Port: 42495 (42495)
- Port Control Protocol
 - Version: 1
 - R bit: Response (1)
 - Opcode: Unknown (0x81)
 - Reserved: 0
 - Result Code: 0
 - Lifetime: 3000
 - Epoch Time: 1576
 - Reserved: 0000000000000000
 - MAP Response
 - Protocol: TCP (6)
 - Reserved: 0
 - Internal Port: 8080
 - Assigned External Port: 8080
 - Assigned External IP Address: ::ffff:161.105.194.14 (::ffff:161.105.194.14)
 - Description Option
 - Option Code: description (0x40)
 - Reserved: 0
 - Option Length: 8 bytes
 - Data: 6965746638350000

The packet bytes pane shows the raw data of the PCP response, including the MAP and Description Option fields. The data for the Description Option is 6965746638350000, which corresponds to the ASCII string "ietf85".

Test#2: SIP-based Communications behind NAT64

SIP Testbed Architecture



PCP & SIP Packets

```
Source: 2001:688:1f94:3000:289f:db7:e8ae:2988 port: 12345
Destination: 2001:688:1f94:3000::2.5351
```

PCP Request:

```
Version: 1
R bit: Request (0)
Opcode: MAP (0x01)
Requested Lifetime: 36000 sec
PCP Client's IP Address: 2001:688:1f94:3000:289f:db7:e8ae:2988
(2001:688:1f94:3000:289f:db7:e8ae:2988)
MAP Request: Protocol: UDP (17)
Internal Port: 3938
Suggested External Port: 3938
Suggested External IP Address: ::ffff:0.0.0.0
Option Code: Unknown (0x7f) Option Length: 12 bytes Data:
0000000000000000000000000000
```

```
Source: 2001:688:1f94:3000::2 port:5153
Destination: 2001:688:1f94:3000:289f:db7:e8ae:2988 port: 12345
```

PCP Response:

```
Version: 1
R bit: Response (1)
Opcode: Unknown (0x81)
Result Code: 0
Lifetime: 36000 sec
Epoch Time: 1
MAP Response Protocol: UDP (17)
Internal Port: 3938
Assigned External Port: 3938
Assigned External IP Address: ::ffff:161.105.194.14 (::ffff:161.105.194.14)
Option Code: PREFIX64 (0x7f) Reserved: 0 Option Length: 12 bytes
Data: 200106881f94300a00000000
```

PCP & SIP Packets

```
REGISTER sip:mysip.fr SIP/2.0
Via: SIP/2.0/UDP 161.105.194.14:3938;branch=z9hG4bK1572043597
From: <sip:client4@mysip.fr:5070>;tag=893886783
To: <sip:client4@mysip.fr:5070>
Call-ID: 1271173454
CSeq: 2 REGISTER
Contact: <sip:client4@161.105.194.14:3938;line=b3433a7df33282d>
    Authorization: Digest username="client4", realm="asterisk",
    nonce="09f75e47", uri="sip:mysip.fr",
    response="826fcff4c6e84ee45fbfa52c351e6316", algorithm=MD5
Max-Forwards: 70
User-Agent: Linphone/3.4.0 (eXosip2/unknown)
Expires: 3600
```

PCP & SIP Packets

```
INVITE sip:13@mysip.fr:5070 SIP/2.0
Via: SIP/2.0/UDP 161.105.194.14:35011;branch=z9hG4bK702695557
From: <sip:client4@mysip.fr:5070>;tag=641336337
To: <sip:13@mysip.fr:5070>
Call-ID: 1532307201
CSeq: 20 INVITE
Contact: <sip:client4@161.105.194.14:35011>
Content-Type: application/sdp
Allow: INVITE, ACK, CANCEL, OPTIONS, BYE, REFER, NOTIFY, MESSAGE, SUBSCRIBE, INFO
Max-Forwards: 70
User-Agent: Linphone/3.4.0 (eXosip2/unknown)
Subject: Phone call
Content-Length: 538
v=0
o=client4 3867 3867 IN IP4 161.105.194.14
s=Talk
c=IN IP4 161.105.194.14
b=AS:256
t=0 0
m=audio 7056 RTP/AVP 111 110 3 101
a=rtpmap:111 speex/16000
a=fmtp:111 vbr=on
a=rtpmap:110 speex/8000
a=fmtp:110 vbr=on
a=rtpmap:3 GSM/8000
a=rtpmap:101 telephone-event/8000
a=fmtp:101 0-11
m=video 9056 RTP/AVP 102 99
a=rtpmap:102 H264/90000
a=fmtp:102 profile-level-id=428014
a=rtpmap:99 MP4V-ES/90000
a=fmtp:99 profile-level-id=3
a=altc: IP6 2001:688:1f94:3000:6c73:ea54:cef:2730 45678
a=altc: IP4 161.105.194.14 7056
```

Encountered Problems

- Android handset disconnects WiFi when no IPv4 address available (no support of DHCPv6)
 - workaround: configure static IPv4 address, wait for RA then delete IPv4
- Address selection for PCP packets

Encountered Problems



The screenshot shows a terminal window titled "Fenêtre 1" with a blue background. The terminal output displays network configuration details for various interfaces. The configuration includes IPv6 addresses and scopes for wlan0, loopback, and host interfaces, along with their respective MTUs, QoS disciplines, and states.

```

inet6 fe80::9a0c:82ff:fec9:c2d0/64 scope link
    valid_lft forever preferred_lft forever
app_112@android:/ # netcfg wlan0 dhcp
^C
130|app_112@android:/ # ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 16436 qdisc noqueue
    state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast
    state UP qlen 1000
    link/ether 98:0c:82:c9:c2:d0 brd ff:ff:ff:ff:ff:ff
3: wlan0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500
    qdisc pfifo_fast state UP qlen 1000
    link/ether 98:0c:82:c9:c2:d0 brd ff:ff:ff:ff:ff:ff
4: rmnet0: <POINTOPOINT,MULTICAST,NOARP> mtu 1500 qdisc noop
    state DOWN qlen 1000
    link/ppp
5: rmnet1: <POINTOPOINT,MULTICAST,NOARP> mtu 1500 qdisc noop
    state DOWN qlen 1000
    link/ppp
6: rmnet2: <POINTOPOINT,MULTICAST,NOARP> mtu 1500 qdisc noop
    state DOWN qlen 1000
    link/ppp
7: ip6tnl0: <NOARP> mtu 1452 qdisc noop state DOWN
    link/tunnel6 :: brd ::
8: sit0: <NOARP> mtu 1480 qdisc noop state DOWN
    link/sit 0.0.0.0 brd 0.0.0.0
9: inet6 ::1/128 scope host
    valid_lft forever preferred_lft forever
10: inet 127.0.0.1/8 scope host lo
    valid_lft forever preferred_lft forever
11: app_112@android:/ #
  
```

Encountered Problems

Packet List:

No.	Time	Source	Destination	Protocol	Length	Info
8	3.098843000	Cisco_d0:1b:8f		DTP	90	Dynamic Trunking Protocol
9	4.154174000	2001:df8:0:112:21bc:8e5b:bc85:ed42	2001:688:1f94:3000::2	PCP	122	Source port: 35578 Destination port: nat-pmp
10	4.154382000	2001:688:1f94:3000::2	2001:df8:0:112:21bc:8e5b:bc85:ed42	PCP	122	Source port: nat-pmp Destination port: 35578

Packet Details:

- Frame 9: 122 bytes on wire (976 bits), 122 bytes captured (976 bits) on interface 0
- Ethernet II, Src: Cisco_d0:1b:c9 (00:1b:90:d0:1b:c9), Dst: 0-link:f4:43:19 (00:0f:3d:f4:43:19)
- Internet Protocol Version 6, Src: 2001:df8:0:112:21bc:8e5b:bc85:ed42, Dst: 2001:688:1f94:3000::2 (2001:688:1f94:3000::2)
- User Datagram Protocol, Src Port: 35578 (35578), Dst Port: nat-pmp (5351)
- Port Control Protocol
 - Version: 1
 - R bit: Request (0)
 - Opcode: MAP (0x01)
 - Reserved: 0
 - Requested Lifetime: 300 sec
 - PCP Client's IP Address: 2001:df8:0:112:9a0c:82ff:fec9:c2d0 (2001:df8:0:112:9a0c:82ff:fec9:c2d0)
 - MAP Request
 - Protocol: TCP (6)
 - Reserved: 0
 - Internal Port: 8080
 - Suggested External Port: 0
 - Suggested External IP Address: ::ffff:0.0.0.0 (::ffff:0.0.0.0)
 - Description Option
 - Option Code: description (0x40)
 - Reserved: 0
 - Option Length: 8 bytes
 - Data: 6965746638350000

Packet Bytes:

```
0000 00 0f 3d f4 43 19 00 1b 90 d0 1b c9 86 dd 60 00  ..=.C.....
0010 00 00 00 44 11 ec 20 01 0d f8 00 00 01 12 21 bc  ...D.....
0020 8e 5b bc 85 ed 42 20 01 06 88 1f 94 30 00 00 00  [...B...0...
0030 00 00 00 00 00 02 8a fa 14 e7 00 44 d5 01 01 01  ....D...
0040 00 00 00 00 01 2c 20 01 0d f8 00 00 01 12 9a 0c  ....
0050 82 ff fe c9 c2 d0 06 00 00 00 1f 90 00 00 00 00  ..ietf85...
0060 00 00 00 00 00 00 00 00 ff ff 00 00 00 00 40 00
0070 00 08 69 65 74 66 38 35 00 00
```

Encountered Problems

- Address mismatch
 - Implementation issue
- Address Lifetime
 - Network APIs does not provide access to lifetime associated to addresses
 - Problem well documented: see RFC6250 and draft-carpenter-behave-referral-object-01
 - pcp-base discusses already considerations related to the refresh of PCP mapping when privacy addresses are in use (section 8.4)
 - Does the base spec need to cite RFC6250?