PCP Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 16, 2013

M. Boucadair
France Telecom

February 12, 2013

Learn NAT64 PREFIX64s using PCP
draft-ietf-pcp-nat64-prefix64-00

Abstract

This document defines a new PCP extension to learn the IPv6 prefix(es) used by a PCP-controlled NAT64 device to build IPv4-embedded IPv6 addresses. This extension is needed for successful communications when IPv4 addresses are used in referrals.

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 August 16, 2013.

Table of Contents

1. Introduction

This document defines a new PCP extension [I-D.ietf-pcp-base] to inform PCP Clients about the Pref64::/n [RFC6052] used by a PCP-controlled NAT64 device [RFC6146].

This extension is required to help establishing communications between IPv6-only hosts and remote IPv4-only hosts.

2. 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 [RFC2119].

3. Problem Statement

3.1. Issues

This document proposes a deterministic solution to solve the following issues:

- Learn the Pref64::/n used by an upstream NAT64 function. This is needed to help:
  - distinguishing between IPv4-converted IPv6 addresses and native IPv6 addresses.
  - implementing IPv6 address synthesis for applications not relying on DNS.
- Avoid stale Pref64::/n.
- Discover multiple Pref64::/n when multiple prefixes in a network.
- Use DNSSEC in the presence of NAT64.

Section 3.2 lists some applications which encounter the issues listed above.

3.2. Use Cases

3.2.1. AAAA Synthesis by Stub-resolver

The extension defined in this document can be used for hosts with DNS64 capability [RFC6147], added to the host's stub-resolver.

The stub resolver on the host will try to obtain (native) AAAA records and if it they are not found, the DNS64 function on the host will query for A records and then synthesizes AAAA records. Using the PREFIX64 PCP extension, the host's stub-resolver can learn the prefix used for IPv6/IPv4 translator and synthesize AAAA records

accordingly.

Learning the Pref64::/n used to construct IPv4-converted IPv6 addresses [RFC6052] allows to make use of DNSSEC.

3.2.2. Applications Referrals

This PCP extension can be used by applications making use of address referrals.

As Peer-to-Peer (P2P) communications for real-time communication is becoming popular with RTCWEB (e.g., P2P for Media, data channels for file transfer etc), this extension can be used to help for NAT64 traversal. SIP [RFC3261] is only one example among those protocols.

3.3. Illustration Example

An illustration example is shown in Figure 1. In this example, NAT64 is co-located with a PCP server while IPv6-only SIP UA interacts with a PCP Client.

In Figure 1, the PCP Client issues a PCP MAP request with PORT_RESERVATION_OPTION to reserve a pair of ports preserving parity and contiguity [I-D.boucadair-pcp-rtp-rtcp]. A pair of ports and an external IPv4 address are then returned by the PCP server to the requesting PCP Client. This information is used by the IPv6-only SIP UA to build its SDP offer which contains exclusively IPv4 addresses (especially in the "c=" line, the port indicated for media port is the external port assigned by the PCP server). The INVITE request including the SDP offer is then forwarded by the NAT64 to the Proxy Server which will relay it to the called party (i.e., IPv4-only SIP UA) (Steps (1) to (3)). IPv4-only SIP UA accepts the offer and sends back its SDP answer in a "200 OK" message which is relayed by the SIP Proxy Server and NAT64 until being delivered to IPv6-only SIP UA (Steps (4) to (6)).

At the end of this process, IPv4-only SIP UA can send media streams to the IPv4 address/port as indicated in the SDP offer while IPv6-only SIP UA can not send media streams as only IPv4 addresses are present in the SDP answer.

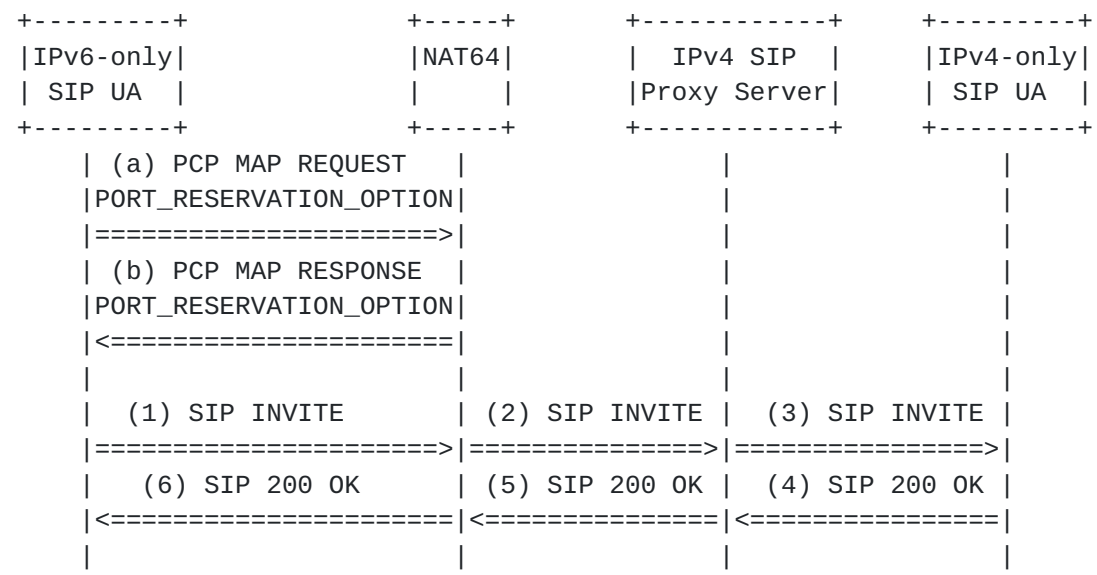

Figure 1

4. PREFIX64 Option

4.1. Format

The format of PREFIX64 PCP Option is depicted in Figure 2.

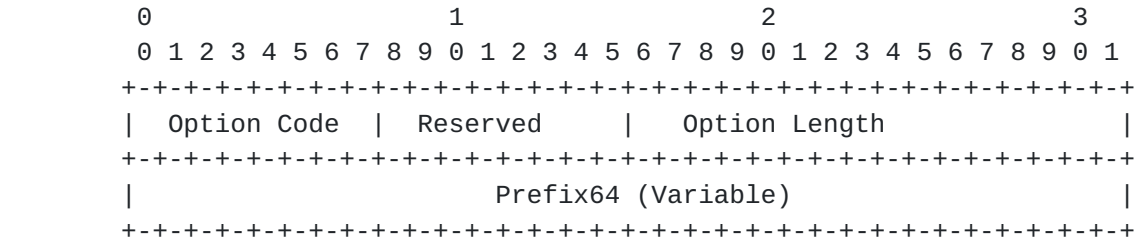

Figure 2: Prefix64 PCP Option

The description of the fields is as follows:

- Option Code: To be assigned by IANA.
- Option Length: Indicates in octets the length of the Pref64::/n. Allowed values are 4, 5, 6, 7, 8, or 12 [RFC6052].
- Prefix64: This field identifies the IPv6 unicast prefix to be used for constructing an IPv4-embedded IPv6 address from an IPv4 address. The address synthesize MUST follow the guidelines documented in [RFC6052].

Option Name: PREFIX64
Number: To be assigned by IANA.
Purpose: Learn the prefix used by the NAT64 to build
IPv4-embedded IPv6 addresses. This is be used by a host
for local address synthesis (e.g., when IPv4 address is
present in referrals).
Valid for Opcodes: MAP
Length: Variable
May appear in: request, response.
Maximum occurrences: 1

4.2. Behaviour

A PCP Client MAY include a PREFIX64 PCP Option in a MAP request to learn the IPv6 prefix used by an upstream PCP-controlled NAT64 device. When enclosed in a MAP request, PREFIX64 MUST be set to ::/96. PREFIX64 PCP Option can be inserted in a MAP request used to learn the external IP address as detailed in Section 11.6 of [I-D.ietf-pcp-base].

A PCP Server controlling a NAT64 SHOULD be configured to return to requesting PCP Clients the value of the Pref64::/n used to build IPv4-embedded IPv6 addresses. When enabled, PREFIX64 PCP Option conveys the value of Pref64::/n.

A PCP Server controlling a NAT64 MAY be configured to inject a PREFIX64 PCP Option in all MAP responses even if the option is not listed in the associated request.

Upon receipt of the PREFIX64 PCP Option, the host embedding the PCP Client uses Pref64::/n for local address synthesize [RFC6052]. How the content of PREFIX64 PCP Option is passed to the OS is implementation-specific.

A PCP Client SHOULD associate each received Pref64::/n with the PCP Server from which the Pref64::/n information was retrieved.

5. Flow Example

Figure 3 shows an example of the use of the option defined in Section 4.

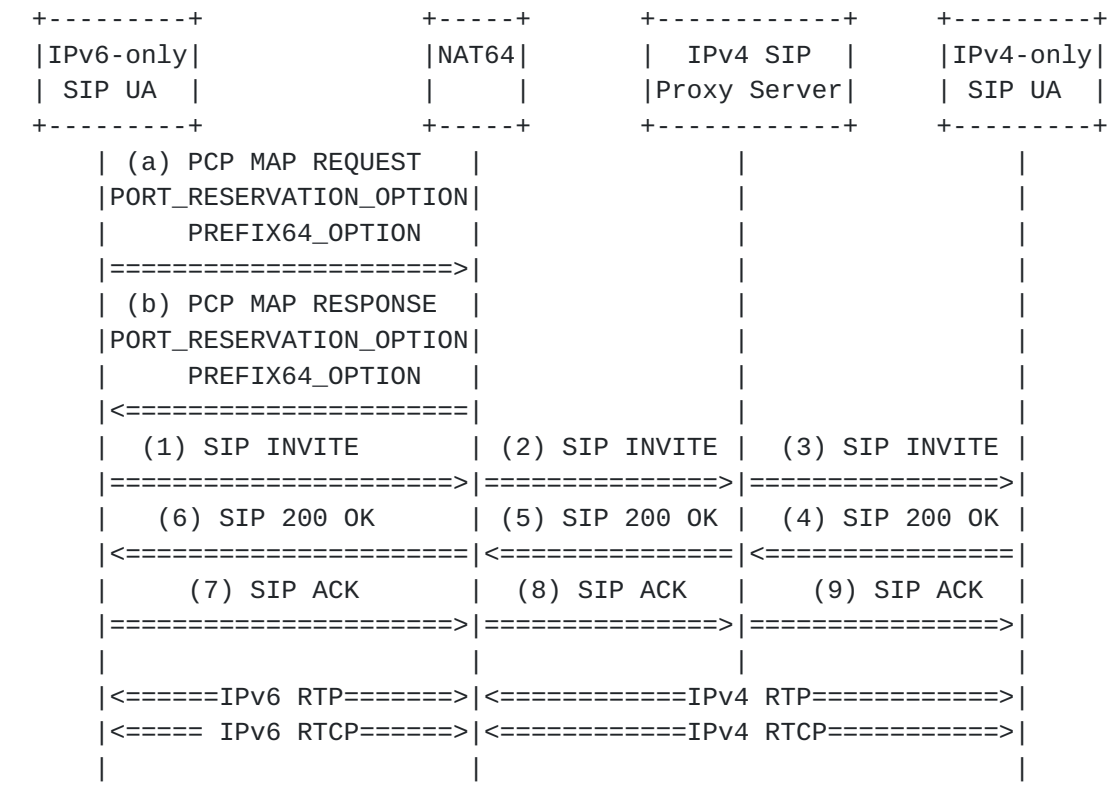

Figure 3: Example of IPv6 to IPv4 SIP initiated Session

In Steps (a) and (b), the IPv6-only SIP UA retrieves a pair of ports to be used for RTP/RTCP, the external IPv4 address and the Pref64::/n to be used to build IPv4-embedded IPv6 addresses. The retrieved IPv4 address and port numbers are used to build the SDP offer in Step (1) while Pref64::/n is used to construct a corresponding IPv6 address of the IPv4 address enclosed in the SDP answer made by the IPv4-only SIP UA (Step 6). RTP/RTCP flows are exchanged between an IPv6-only SIP UA and an IPv4-only UA without requiring any ALG at the NAT64 and no particular function to be supported by the IPv4-only SIP Proxy Server to help establishing the session (e.g., Hosted NAT traversal).

When the session is initiated from IPv4 SIP UA (see Figure 4): Steps (a) and (b), the IPv6-only SIP UA retrieves a pair of ports to be used for RTP/RTCP, the external IPv4 address and the Pref64::/n to be used to build IPv4-embedded IPv6 addresses. These two steps can be delayed until receiving the INVITE message (Step 3).

The retrieved IPv4 address and port numbers are used to build the SDP answer in Step (4) while Pref64::/n is used to construct a corresponding IPv6 address of the IPv4 address enclosed in the SDP offer made by the IPv4-only SIP UA (Step 3). RTP/RTCP flows are exchanged between an IPv6-only SIP UA and an IPv4-only UA without

requiring any ALG at the NAT64 and no particular function to be supported by the IPv4-only SIP Proxy Server to help establishing the session (e.g., Hosted NAT traversal).

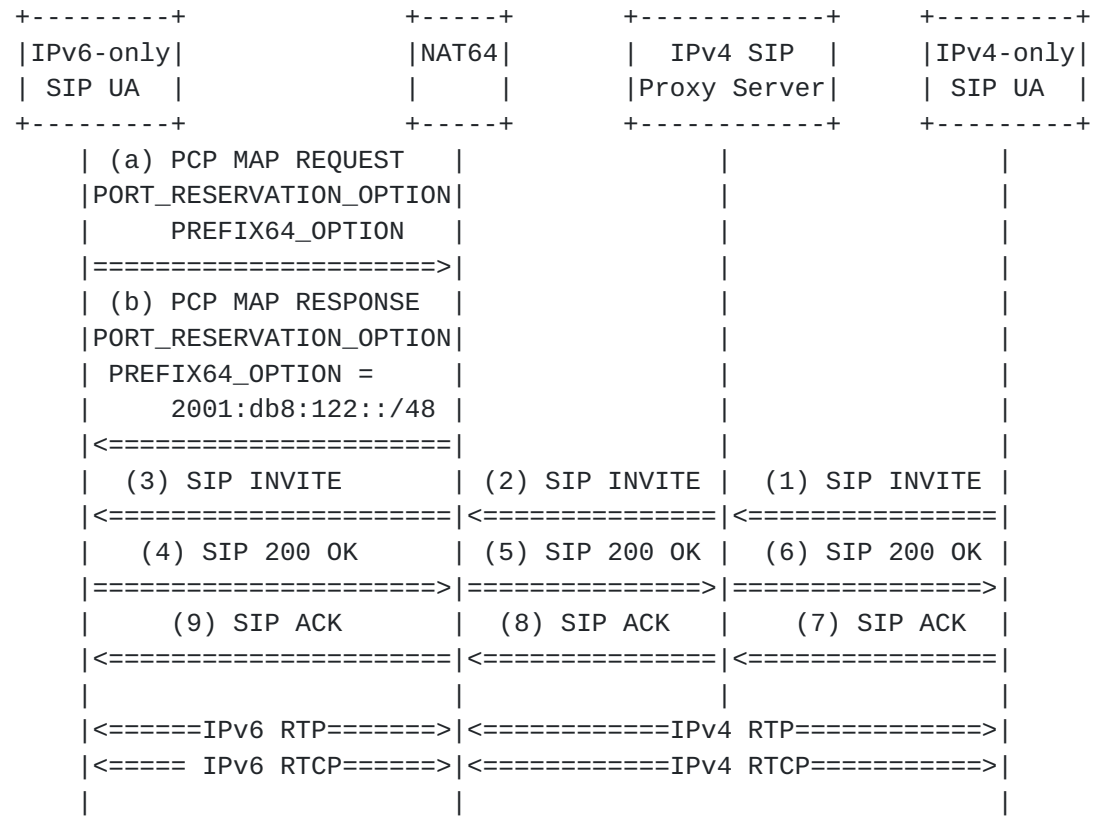

Figure 4: Example of IPv4 to IPv6 SIP initiated Session

6. IANA Considerations

This document requests a new PCP option:
PREFIX64

7. Security Considerations

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

8. Acknowledgements

Many thanks to S. Perreault , R. Tirumaleswar, T. Tsou, D. Wing, J.

Zhao and R. Penno for the comments and suggestions.

9. References

9.1. 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-29 (work in progress), November 2012.

[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.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, October 2010.

[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.

[RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, April 2011.

9.2. Informative References

[I-D.boucadair-pcp-rtp-rtcp]
Boucadair, M. and S. Sivakumar, "Reserving N and N+1 Ports with PCP", draft-boucadair-pcp-rtp-rtcp-05 (work in progress), October 2012.

Author's Address

Mohamed Boucadair
France Telecom
Rennes, 35000
France

Email: mohamed.boucadair@orange.com