Public IPv4 over IPv6 Access Network
draft-ietf-softwire-public-4over6-09

Abstract

   When the service provider networks are upgraded to IPv6, end users
   will continue to demand IPv4 connectivity. This document describes a
   mechanism for hosts or customer networks in IPv6 access network to
   build bidirectional IPv4 communication with the IPv4 Internet. The
   mechanism follows the hub and spokes softwire model, and uses IPv4-
   over-IPv6 tunnel as basic method to traverse IPv6 network. The bi-
   directionality of this IPv4 communication is achieved by explicitly
   allocating public IPv4 addresses to end users, as well as maintaining
   IPv4-IPv6 address binding on the border relay. This mechanism
   features the allocation of full IPv4 address over IPv6 network, and
   has been used in production for high-end IPv4 users, IPv6 transition
   of ICPs, etc.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

   When deploying IPv6 networks, IPv4 connectivity is still a
   functionality required by end users. It is used for IPv4
   communication with IPv4-only part of the Internet during the IPv4-
   IPv6 transition period. IPv4-over-IPv6 tunnel mechanisms are the
   general solutions to provide this type of IPv4 services.

   This document describes a mechanism for providing IPv4 connectivity
   in this situation. The mechanism is similar to the Binding approach
   of the Unified IPv4-in-IPv6 Softwire CPE effort that is documented in
   [I-D.bfmk-softwire-unified-cpe] Section 2. Although the
   functionality documented in Unified CPE is similar, this document
   describes existing practice that differs from Unified CPE, but that
   has been deployed in China Next Generation Internet (CNGI) - China
   Education and Research Network 2 (CERNET2).

   The purpose of this draft is to document the protocol that was
   deployed, both for historical purposes and for the benefit of users
   of that protocol in the field at the time of publication. Future
   deployments with similar requirements should simply use the related
   mechanism in [I-D.bfmk-softwire-unified-cpe].

   The advantage of IPv4-over-IPv6 tunnel mechanisms is the transparency
   to the IPv6 infrastructure: since IPv4 is actually only needed on the
   end user side as well as beyond the tunnel concentrator, most
   networks and services within the ISP network can remain IPv6 only.
   Therefore, operators can run an IPv6-only infrastructure instead of a
   fully dual-stack network, as well as save the IPv4 address resource
   from being assigned all over the network.

   While different IPv4-over-IPv6 mechanisms are developed for different
   application scenarios, the mechanism described in this document
   focuses on providing full end-to-end transparency to the user-side.
   Therefore, public IPv4 addresses are expected to be provisioned to
   end users and carrier-side address translation can be avoided.
   Further more, full address is preferred to port-restricted address.
   With full address provisioned, user-side address translation is not
   necessarily needed either. This means minimal changes to the user
   side: operating system could support the mechanism smoothly, while
   transparency on upper-layer applications is guaranteed. For many
   ISPs which are actually capable of provisioning full IPv4 addresses,
   the mechanism provide a pure, suitable solution.

   Another focus of this mechanism is deployment and operation
   flexibility. This mechanism allows IPv4 addressing and IPv6
   addressing schemes to be independent of each other: end user IPv4
   address is not embedded in its IPv6 address. The IPv6 infrastructure

in the middle is not involved with the IPv4-over-IPv6 mechanism, so
no special network planning is required; the service can be provided
in on-demand style; the IPv4 address resources can be managed in a
flat, centralized manner rather than in the way of distributed to
customer sites with IPv6. This requires the border relay to maintain
the binding of IPv4 address and IPv6 address dynamically, i.e.
maintaining per-subscriber binding state.

The mechanism follows hub and spokes softwire model, and uses IPv4-
over-IPv6 tunnel between end host or CPE and border relay as basic
data plane method. Full IPv4 addresses are allocated from the ISP to
the end host or CPE over IPv6 network. Simultaneously, the binding
between the allocated IPv4 address and the end user's IPv6 address
are maintained on the border relay for encapsulation usage.

2. Terminology

Public 4over6: Public 4over6 is a per-subscriber stateful, IPv4-over-
IPv6 tunnel mechanism proposed by this document. Public 4over6
supports bidirectional communication between IPv4 Internet and IPv4
hosts or customer networks in IPv6 access network, by leveraging
IPv4-in-IPv6 tunnel and public IPv4 address allocation over IPv6.

4over6 Customer Edge (CE): A device functioning as a Customer Edge
equipment in Public 4over6 environment. The 4over6 CE can be either
a dual-stack capable host, or a dual-stack CPE device, both of which
have a tunnel interface to support IPv4-in-IPv6 encapsulation. In
the former case, the host supports both IPv4 and IPv6 stack but is
provisioned with IPv6 only. In the latter case, the CPE has an IPv6
interface connecting to ISP network, and an IPv4 or dual-stack
interface connecting to customer network; hosts in the customer
network can be IPv4-only or dual-stack.

4over6 Border Relay (BR): A router functioning as the border relay in
Public 4over6 environment. 4over6 BR is the IPv4-in-IPv6 tunnel
concentrator located in IPv6 ISP network. It is a dual-stack router
which connects to both the IPv6 ISP network and IPv4 Internet. The
4over6 BR also works as a DHCPv4 over IPv6
[I-D.ietf-dhc-dhcpv4-over-ipv6] server/relay for assigning public
IPv4 address to 4over6 CEs.

3. Scenario and Use Cases

The general scenario of Public 4over6 is shown in Figure 1. Users in
an IPv6 network take IPv6 as their native service. Some users are
end hosts which face the ISP network directly, while the others are
customer LAN networks behind CPEs, such as a home LAN, an enterprise
network, etc. The ISP network is IPv6-only rather than dual-stack,

which means the ISP cannot provide native IPv4 service to users. In
order to support legacy IPv4 transport, some routers on the carrier
side are dual-stack and are connected to IPv4 Internet. These
routers act as 4over6 border relays. Network users that require IPv4
connectivity obtain it using these routers.

```
                  +-------------------------+
                  |    IPv6 ISP Network      |
             +------+                        |
             |4over6|Host                    |
             |  CE  |================+-------+   +-----------+
             +------+                |4over6 |   |   IPv4    |
 +---------+     |      IPv4-in-IPv6 |  BR   |---| Internet  |
 |Customer |  +------+               |       |   |           |
 |IPv4 LAN |--|4over6|================+-------+   +-----------+
 | Network |  |  CE  |CPE                   |
 +---------+  +------+                      |
                  |                         |
                  +-------------------------+
```
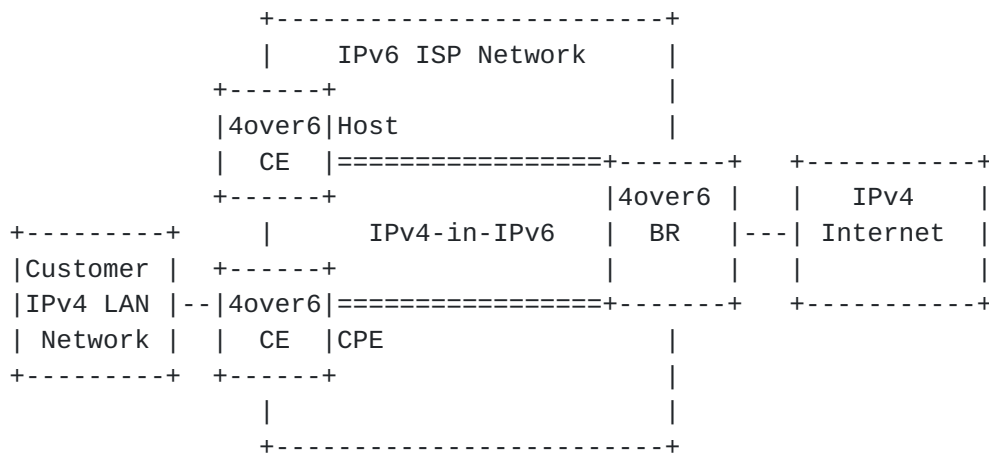
Figure 1 Public 4over6 scenario

Public 4over6 can be applicable in several use cases. If an ISP
which switches to IPv6 still has plenty of IPv4 address resource, it
can deploy Public 4over6 to provide transparent IPv4 service for all
its customers. If the ISP does not have so much IPv4 addresses, it
can deploy Dual-Stack Lite [RFC6333] as the basic IPv4-over-IPv6
service. Along with DS-Lite, Public 4over6 can be deployed as a
value-added service, overcoming the service degradation caused by the
CGN. The two mechanisms can be integrated, because the IPv4-in-IPv6
tunnel functions are the same; the difference is that DS-Lite employs
a CGN while Public 4over6 employs an IPv4 provisioning process. A
typical case of the high-end users that could use Public 4over6 is
IPv4 application server. Full, public IPv4 address brings
significant convenience in this case, which is important to IPv6
transition for ICPs. The DNS registration can be direct using
dedicated address; the access of the application service can be
straightforward, with no translation involved; there will be no need
to provide NAT traversal mechanisms for incoming traffic, and no
special handling is required for well-known ports.

4. Public 4over6 Address Provisioning

4.1. Basic Provisioning Steps

The following figure shows the basic provisioning steps for Public
4over6.


```
          4over6                    DHCPv6          4over6         DHCPv4
           CE                       Server            BR           Server
            |Assign IPv6 Addr/Pref +|                |              |
            |  BR's IPv6 Addr Info  |                |              |
            |<----------------------|                |              |
            |       DHCPv6/Other    |                |              |
          WAN                                        |              |
        IPv6 Configure                               |              |
            |                                        |              |
            | Assign Public IPv4 Addr(DHCPv4-over-v6/Static Conf) |
            |<-------------------------------------|<-------------|
            |                                        | IPv4-IPv6    |
            |                                        | Binding SYN  |
          Tunnel                                     |
        IPv4 Configure                        Binding Update
            |                                        |
            |            IPv4-in-IPv6 Tunnel         |
            |<-------------------------------------->|
            |                                        |
```


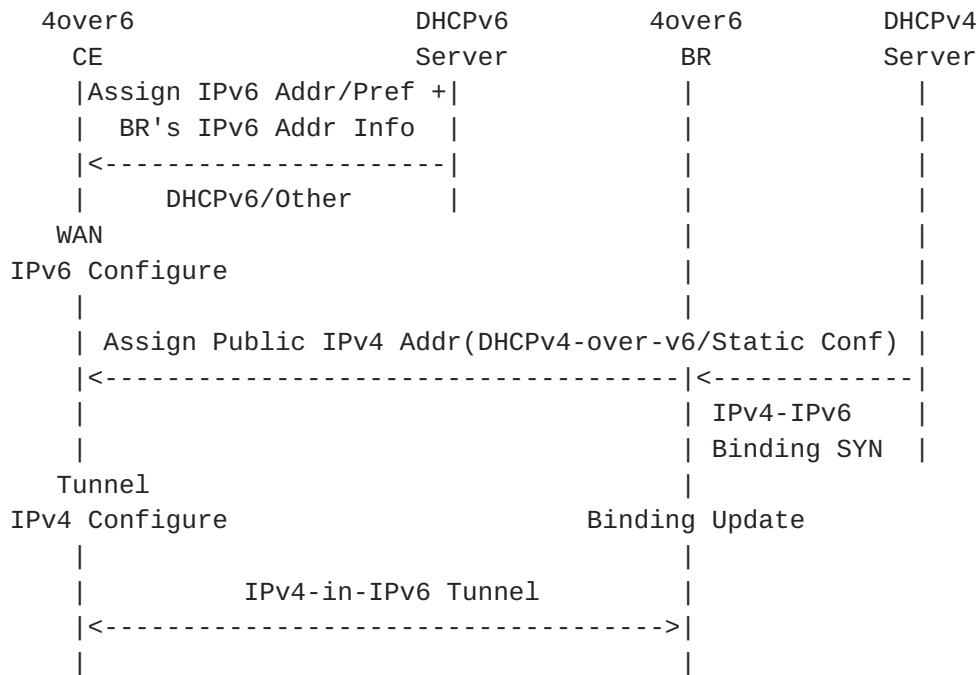   Figure 2 Public 4over6 Address Provisioning

   The main steps are:

   o Provision IPv6 address/prefix to 4over6 CE, along with the
      information of 4over6 BR's IPv6 address, by DHCPv6 or other means.

   o 4over6 CE configures its WAN interface with globally unique IPv6
      address which is a result of IPv6 provisioning, including DHCPv6,
      SLAAC or manual configuration.

   o Provision IPv4 address to 4over6 CE, by DHCPv4 over IPv6 or static
      configuration.

   o 4over6 BR obtains the IPv4 and IPv6 addresses of the 4over6 CE
      using information provided by the DHCPv4 sever.

   o 4over6 CE configures its tunnel interface, as a result of IPv4
      provisioning.

   o 4over6 BR updates the IPv4-IPv6 address binding table, as a result
      of address binding information acquired from DHCPv4 server.

4.2. Public IPv4 Address Allocation

Usually each CE is provisioned with one public IPv4 address. However
it is possible that a CE would require an IPv4 prefix. The key
problem here is the mechanism for IPv4 address provisioning over IPv6
network.

There are two possibilities: DHCPv4 over IPv6, and static
configuration. Public 4over6 supports both these methods. DHCPv4
over IPv6 enables DHCPv4 message to be transported in IPv6 rather
than IPv4; therefore, the DHCPv4 process can be performed over an
IPv6 network, between BR and CE. [I-D.ietf-dhc-dhcpv4-over-ipv6]
describes the DHCP protocol extensions to support that. As to static
configuration, 4over6 users and the ISP operators negotiate
beforehand to authorize the IPv4 address(es). Then the tunnel
interface and the address binding are configured by the user and the
ISP respectively.

While regular users would probably take DHCPv4 over IPv6, the manual
configuration is usually seen in two cases: application server, which
requires a stable IPv4 address, and enterprise network, which usually
requires an IPv4 prefix rather than one single address (Note that
DHCPv4 does not support prefix allocation).

5. 4over6 CE Behavior

A CE is provisioned with IPv6 before Public 4over6. It also learns
the BR's IPv6 address beforehand. This IPv6 address can be
configured using a variety of methods, ranging from an out-of-band
mechanism, manual configuration, or DHCPv6 option. In order to
guarantee interoperability, the CE element implements the AFTR-Name
DHCPv6 option defined in [RFC6334].

A CE supports DHCPv4 over IPv6[I-D.ietf-dhc-dhcpv4-over-ipv6], to
dynamically acquire IPv4 address over IPv6 and assign it to the IPv4-
in-IPv6 tunnel interface. The CE regards the BR as DHCPv4-over-IPv6
server/relay for public IPv4 address allocation, whose IPv6 address
is learned by the CE as described above.

A CE also supports static configuration of the tunnel interface. In
the case of prefix provisioning, the tunnel interface is assigned
with the Well-Known IPv4 Address defined in section 5.7 of [RFC6333],
rather than using an address from the prefix. If the CE has multiple
IPv6 addresses on its WAN interface, it uses the same IPv6 address
for DHCPv4 over IPv6/negotiation of manual configuration, and for
data plane encapsulation.

A CE performs IPv4-in-IPv6 encapsulation and decapsulation on the

tunnel interface. When sending out an IPv4 packet, it performs the
encapsulation, using the IPv6 address of the 4over6 BR as the IPv6
destination address, and its own IPv6 address as the IPv6 source
address. The decapsulation on 4over6 CE is simple. When receiving
an IPv4-in-IPv6 packet, the CE just removes the IPv6 header, and
either hands it to upper layer or forward it to customer network
according to the IPv4 destination address.

A CE runs a regular IPv4 NAPT for its customer network when it is
provisioned with one single IPv4 address. In that case, the assigned
IPv4 address of the tunnel interface would be the external IPv4
address of the NAPT. Then the CE performs IPv4 private-to-public
translation before encapsulation of IPv4 packets from the customer
network, and IPv4 public-to-private translation after decapsulation
of IPv4-in-IPv6 packets.

IPv4 NAPT is not necessarily when the CE is provisioned with an IPv4
prefix. In this case, the detailed customer network planning is out
of scope.

4over6 CE supports backward compatibility with DS-Lite. A CE can
employ Well-Known IPv4 Address for B4 [RFC6333] and switch to Dual-
Stack Lite for IPv4 communications, if it can't get a public IPv4
address from the DHCPv4 server (for instance, the DHCPv4 over IPv6
process fails or the DHCPv4 server refuses to allocate a public IPv4
address to it, etc.).

6. 4over6 BR Behavior

4over6 BR maintains the bindings between the CE IPv6 address and CE
IPv4 address (prefixes). The bindings are used to provide correct
encapsulation destination address for inbound IPv4 packets, as well
as validate the IPv6-IPv4 source of the outbound IPv4-in-IPv6
packets.

The BR acquires the binding information through the IPv4 address
provisioning process. For static configuration, the operator
manually configures the BR using the binding information obtained
through negotiation with the customer. As for DHCPv4-over-IPv6,
there are multiple possibilities which are deployment-specific:

o The BR can be collocated with the DHCPv4-over-IPv6 server. Then
   the synchronization happens within the BR. It installs a binding
   when sending out an ACK for a DHCP lease, and delete it when the
   lease expires or a DHCP RELEASE is received.

o The BR can play the role of TRA as described in
   [I-D.ietf-dhc-dhcpv4-over-ipv6], and snoop for the DHCPv4 ACK and

Release messages, as well as keep a timer for each binding
according to the DHCP lease time.

On the IPv6 side, the BR decapsulates IPv4-in-IPv6 packets coming
from 4over6 CEs. It removes the IPv6 header of every IPv4-in-IPv6
packet and forwards it to the IPv4 Internet. Before the
decapsulation, the BR checks the inner IPv4 source address against
the outer IPv6 source address, by matching such a binding entry in
the binding table. If no binding is found, the BR silently drops the
packet. On the IPv4 side, the BR encapsulates the IPv4 packets
destined to 4over6 CEs. When performing the IPv4-in-IPv6
encapsulation, the BR uses its own IPv6 address as the IPv6 source
address, uses the IPv4 destination address in the packet to look up
IPv6 destination address in the address binding table. After the
encapsulation, the BR sends the IPv6 packet on its IPv6 interface to
reach a CE.

The BR supports hairpinning of traffic between two CEs, by performing
de-capsulation and re-encapsulation of packets.

7. Fragmentation and reassembly

The same considerations as described in section 5.3 and section 6.3
of [RFC6333] are taken into account.

8. DNS

The procedure described in Section 5.5 and Section 6.4 of [RFC6333]
is followed.

9. IANA Considerations

This document has no IANA actions.

10. Security Considerations

In some cases, the 4over6 BR is expected to implement methods to
limit service only to registered customers. One simple solution is
to make sure the BR allocates IPv4 addresses only to registered
customers, i.e. the BR can filter on the IPv6 source addresses of
incoming DHCP packets and only respond to the ones which have
registered IPv6 source address. The BR can also perform validation
during DHCP, for example, based on the MAC address of the CEs. As to
data packets, the BR can implement an IPv6 ingress filter on the
tunnel interface to accept only the IPv6 address range defined in the
filter, as well as check the IPv4-IPv6 source address binding by
looking up the binding table.

11. **Change Log from the -03 Version (RFC Editors please remove this part)**

   1. Change the Intended Status to Informational, and reword some text to not use RFC2119 language.

   2. Specify the feature of Public 4over6 and circumstances requiring the mechanism in Abstract.

   3. Explain the motivation of IPv4-over-IPv6 for Public 4over6 in section 1.

   4. Explain the relationship between Public 4over6 and Unified CPE, as well as the purpose of this doc.

   5. Clarify that customer network behind the 4over6 CE could be IPv4-only or dual-stack in section 3.

   6. Explain how to integrate Public 4over6 and DS-lite as a typical use case in section 4 and section 5.

   7. Clarify that IPv6 address/prefix can both be supported by 4over6 CEs in section 5.

   8. Improve the preciseness of the texts.

   9. Remove the text that describes the BR not participating the DHCPv4-over-IPv6 process.

   10. Update the references.

12. **Author List**

   The following are extended authors who contribute to the effort:

      Huiling Zhao
      China Telecom
      Room 502, No.118, Xizhimennei Street
      Beijing 100035
      P.R.China

      Phone: +86-10-58552002
      Email: zhaohl@ctbri.com.cn

      Chongfeng Xie
      China Telecom

Room 708, No.118, Xizhimennei Street
Beijing 100035
P.R.China

Phone: +86-10-58552116
Email: xiechf@ctbri.com.cn

Qiong Sun
China Telecom
Room 708, No.118, Xizhimennei Street
Beijing 100035
P.R.China

Phone: +86-10-58552936
Email: sunqiong@ctbri.com.cn

Qi Sun
Tsinghua University
Beijing 100084
P.R.China

Phone: +86-10-62785822
Email: sunqi@csnet1.cs.tsinghua.edu.cn

Chris Metz
Cisco Systems
3700 Cisco Way
San Jose, CA 95134
USA

Email: chmetz@cisco.com

[13](#). References

[13.1](#). Normative References

[RFC4925]   Li, X., Dawkins, S., Ward, D., and
                             A. Durand, "Softwire Problem
                             Statement", RFC 4925, July 2007.

[RFC4966]   Aoun, C. and E. Davies, "Reasons to
                             Move the Network Address Translator

                                    - Protocol Translator (NAT-PT) to
                                    Historic Status", RFC 4966,
                                    July 2007.

   [RFC5549]   Le Faucheur, F. and E. Rosen,
                                    "Advertising IPv4 Network Layer
                                    Reachability Information with an
                                    IPv6 Next Hop", RFC 5549, May 2009.

   [RFC5565]   Wu, J., Cui, Y., Metz, C., and E.
                                    Rosen, "Softwire Mesh Framework",
                                    RFC 5565, June 2009.

   [RFC6333]   Durand, A., Droms, R., Woodyatt, J.,
                                    and Y. Lee, "Dual-Stack Lite
                                    Broadband Deployments Following IPv4
                                    Exhaustion", RFC 6333, August 2011.

   [RFC6334]   Hankins, D. and T. Mrugalski,
                                    "Dynamic Host Configuration Protocol
                                    for IPv6 (DHCPv6) Option for Dual-
                                    Stack Lite", RFC 6334, August 2011.

13.2. Informative References

   [I-D.bfmk-softwire-unified-cpe] Boucadair, M. and I. Farrer,
                                    "Unified IPv4-in-IPv6 Softwire CPE",
                                    draft-bfmk-softwire-unified-cpe-02
                                    (work in progress), January 2013.

   [I-D.ietf-dhc-dhcpv4-over-ipv6] Cui, Y., Wu, P., Wu, J., and T.
                                    Lemon, "DHCPv4 over IPv6 Transport",
                                    draft-ietf-dhc-dhcpv4-over-ipv6-06
                                    (work in progress), March 2013.

Authors' Addresses

   Yong Cui
   Tsinghua University
   Department of Computer Science, Tsinghua University
   Beijing 100084
   P.R.China

   Phone: +86-10-6260-3059
   EMail: yong@csnet1.cs.tsinghua.edu.cn

   Jianping Wu

      Tsinghua University
      Department of Computer Science, Tsinghua University
      Beijing 100084
      P.R.China

      Phone: +86-10-6278-5983
      EMail: jianping@cernet.edu.cn

      Peng Wu
      Tsinghua University
      Department of Computer Science, Tsinghua University
      Beijing 100084
      P.R.China

      Phone: +86-10-6278-5822
      EMail: pengwu.thu@gmail.com

      Olivier Vautrin
      Juniper Networks
      1194 N Mathilda Avenue
      Sunnyvale, CA 94089
      USA

      EMail: Olivier@juniper.net

      Yiu L. Lee
      Comcast
      One Comcast Center
      Philadelphia, PA 19103
      USA

      EMail: yiu_lee@cable.comcast.com