

Internet Engineering Task Force (IETF)
Request for Comments: 7040
Category: Informational
ISSN: 2070-1721

Y. Cui
J. Wu
P. Wu
Tsinghua University
O. Vautrin
Juniper Networks
Y. Lee
Comcast
November 2013

Public IPv4-over-IPv6 Access Network

Abstract

This document describes a mechanism called Public 4over6, which is designed to provide IPv4 Internet connectivity over an IPv6 access network using global IPv4 addresses. Public 4over6 was developed in the IETF and is in use in some existing deployments but is not recommended for new deployments. Future deployments of similar scenarios should use Lightweight 4over6. Public 4over6 follows the Hub and Spoke software model and uses an IPv4-in-IPv6 tunnel to forward IPv4 packets over an IPv6 access network. The bidirectionality of the IPv4 communication is achieved by explicitly allocating global non-shared IPv4 addresses to end users and by maintaining IPv4-IPv6 address binding on the border relay. Public 4over6 aims to provide uninterrupted IPv4 services to users, like Internet Content Providers (ICPs), etc., while an operator makes the access network transition to an IPv6-only access network.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7040>.

Copyright Notice

Copyright (c) 2013 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. Terminology	4
3. Scenario and Use Cases	4
4. Public 4over6 Address Provisioning	6
4.1. Basic Provisioning Steps	6
4.2. Public IPv4 Address Allocation	7
5. 4over6 CE Behavior	7
6. 4over6 BR Behavior	8
7. Fragmentation and Reassembly	9
8. DNS	9
9. Security Considerations	10
10. Contributors	11
11. References	12
11.1. Normative References	12
11.2. Informative References	12

1. Introduction

When operators make the access network transition to an IPv6-only access network, they must continue to provide IPv4 services to their users to access IPv4 contents. IPv4 connectivity is required when communicating with the IPv4-only Internet. This document describes a mechanism called Public 4over6 for providing IPv4 connectivity over a native IPv6-only access network. This memo focuses on interactions between Public 4over6 elements as well as the deployment architecture.

Public 4over6 is in active deployment in some environments, particularly in China Next Generation Internet (CNGI) and China Education and Research Network 2 (CERNET2), but it is not recommended for new deployments. Documenting this approach is intended to benefit users and operators of existing deployments as well as readers of other IPv4-over-IPv6 documents.

In addition to Public 4over6 and its deployment architecture as described in this memo, the IETF is currently working on a more generic solution called Lightweight 4over6 [SOFTWARE-LW46], which is classified as a binding approach in the Unified IPv4-in-IPv6 Software Customer Premises Equipment (CPE) [SOFTWARE-CPE]. Lightweight 4over6 covers both sharing and non-sharing global IPv4 addresses in the Hub and Spoke model. Future deployments should use [SOFTWARE-LW46].

Public 4over6 utilizes the IPv4-in-IPv6 tunnel technique defined in [RFC2473], which enables IPv4 datagrams to traverse through native IPv6 networks. IPv4 nodes connect to the Tunnel Entry-Point Node and Tunnel Exit-Point Node to communicate over the IPv6-only network. Therefore, the Internet Service Providers (ISPs) can run an IPv6-only infrastructure instead of a fully dual-stack network as well as avoid the need to deploy scarce IPv4 address resources throughout the network.

This mechanism focuses on providing full end-to-end transparency to the user side. Therefore, global IPv4 addresses are expected to be provisioned to end users, and carrier-side address translation can be avoided. Furthermore, global non-shared IPv4 addresses are preferable to shared IPv4 addresses, so that user-side address translation is not necessary either. It is important, in particular, to users that are required to run their applications on an IP protocol different from TCP and UDP (e.g., IPsec, L2TP) or on certain well-known TCP/UDP ports (e.g., HTTP, SMTP). For many ISPs that are actually capable of provisioning non-shared unique IPv4 addresses, the mechanism provides a pure, suitable solution.

Another focus of this mechanism is deployment and operational flexibility. Public 4over6 allows IPv4 and IPv6 address architectures to be totally independent of each other; the end user's IPv4 address is not embedded in its IPv6 address. Therefore, IPv4 address planning has no implication for IPv6 address planning. Operators can manage the IPv4 address resources in a flat, centralized manner. This requires that the tunnel concentrator [RFC4925] maintain the binding between an IPv4 address and an IPv6 address, i.e., maintaining per-subscriber binding state.

The mechanism follows the Hub and Spoke software model [RFC4925] and uses IPv4-in-IPv6 tunneling as the basic data-plane method. Global non-shared IPv4 addresses are allocated from the ISP to end hosts or CPEs over an IPv6 network. Simultaneously, the binding between the allocated IPv4 address and the end user's IPv6 address is maintained on the tunnel concentrator for encapsulation usage.

2. Terminology

Public 4over6: A per-subscriber, stateful IPv4-in-IPv6 tunnel mechanism. Public 4over6 supports bidirectional communication between the global IPv4 Internet and IPv4 hosts or customer networks via an IPv6 access network by leveraging IPv4-in-IPv6 tunneling [RFC2473] and global IPv4 address allocation over IPv6. The term 'Public' means the allocated IPv4 address is globally routable.

Full IPv4 address: An IPv4 address that is not shared by multiple users. The user with this IPv4 address has full access to all the available TCP/UDP ports, including the well-known TCP/UDP ports.

4over6 Customer Edge (CE): A device functioning as the Customer Edge equipment in a Public 4over6 environment. A 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 stacks but its uplink is IPv6 only. In the latter case, the CPE has an IPv6 interface connecting to the ISP network and an IPv4 or dual-stack interface connecting to the customer network; hosts in the customer network can be IPv4 only or dual stack.

4over6 Border Relay (BR): A router deployed in the edge of the operator's IPv6 access network that supports IPv4-in-IPv6 tunnel termination. A 4over6 BR is a dual-stack router that connects to both the IPv6 access network and the IPv4 Internet. The 4over6 BR can also work as a DHCPv4-over-IPv6 [DHCPv4-IPv6] server/relay for assigning and distributing global IPv4 addresses to 4over6 CEs.

3. Scenario and Use Cases

The general Public 4over6 scenario is shown in Figure 1. Users in an IPv6 network take IPv6 as their native service. Some users are end hosts that face the ISP network directly, while others are in private networks behind CPEs, such as a home Local Area Network (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 the IPv4 Internet. These routers act as 4over6 BRs. Network users that require IPv4 connectivity obtain it through these routers.

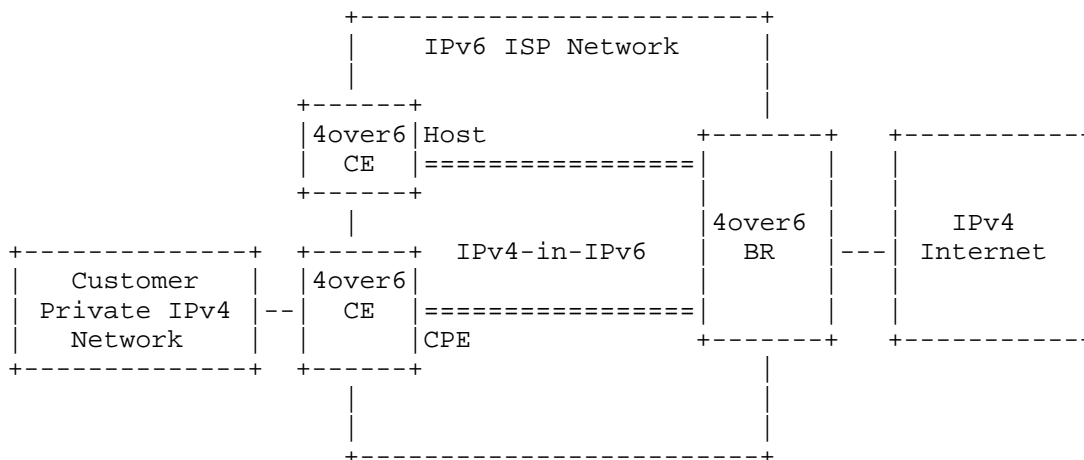


Figure 1: Public 4over6 Scenario

Public 4over6 can be applicable in several use cases. If an ISP that switches to IPv6 still has plenty of global IPv4 address resources, it can deploy Public 4over6 to provide transparent IPv4 service for all its customers. If the ISP does not have enough IPv4 addresses, it can deploy Dual-Stack Lite [RFC6333] as the basic IPv4-over-IPv6 service. Along with Dual-Stack Lite, Public 4over6 can be deployed as a value-added service, overcoming the service degradation caused by the Carrier Grade NAT (CGN). An IPv4 application server is a typical high-end user of Public 4over6. Using a full, global IPv4 address brings significant advantages in this case and is important for Internet Content Providers (ICPs) making the transition to IPv6:

- o The DNS registration can be direct, using a dedicated address;
- o Accessing the application service can be straightforward, with no translation involved;
- o There will be no need to provide NAT traversal mechanisms for incoming traffic, and no special handling is required for the well-known TCP/UDP ports.

4. Public 4over6 Address Provisioning

4.1. Basic Provisioning Steps

Figure 2 shows the basic provisioning steps for Public 4over6.

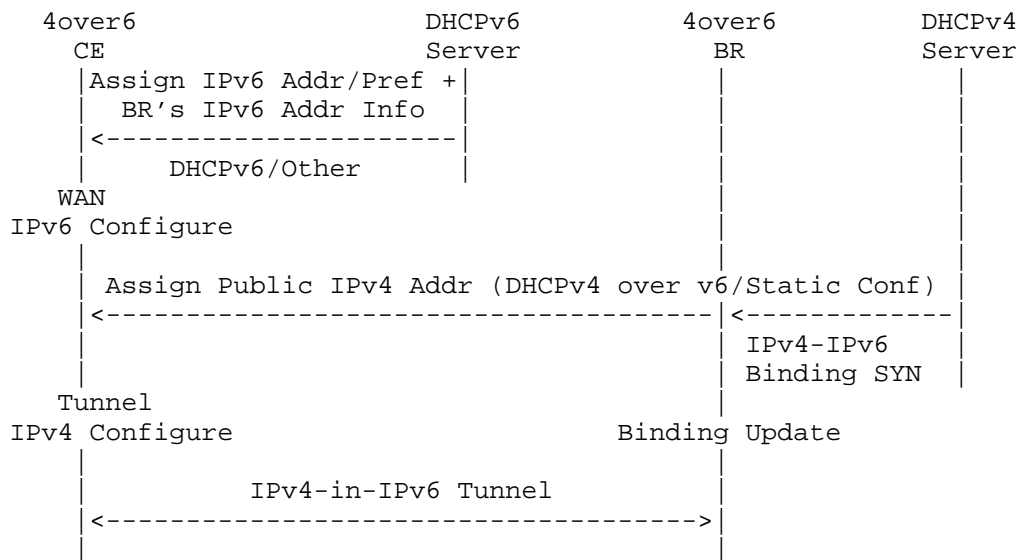


Figure 2: Public 4over6 Address Provisioning

The main steps are:

- o IPv6 address/prefix is provisioned to 4over6 CE, along with knowledge of 4over6 BR's IPv6 address, using DHCPv6 or other means.
- o 4over6 CE configures its WAN interface with a globally unique IPv6 address, which is a result of IPv6 provisioning, including DHCPv6, Stateless Address Autoconfiguration (SLAAC), or manual configuration.
- o IPv4 address is provisioned 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 server.

- o 4over6 CE configures its tunnel interface as a result of IPv4 provisioning.
- o 4over6 BR updates the IPv4-IPv6 address-binding table according to the address-binding information acquired from the DHCPv4 server.

4.2. Public IPv4 Address Allocation

Usually, each CE is provisioned with one global 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 networks.

There are two possibilities: DHCPv4 over IPv6, and static configuration. Public 4over6 supports both these methods. DHCPv4 over IPv6 allows DHCPv4 messages to be transported in IPv6 rather than IPv4; therefore, the DHCPv4 process can be performed over an IPv6 network between the BR and the relevant CE. [DHCPv4-IPv6] describes the DHCP protocol extensions needed to support this operation. For static configuration, Public 4over6 users and 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 opt for DHCPv4 over IPv6, the static configuration is particularly applicable in two cases: for application servers, which require a stable IPv4 address; and for enterprise networks, which usually require 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 the Public 4over6 process. 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 via a 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 [DHCPv4-IPv6] to dynamically acquire an IPv4 address over IPv6 and assign it to the IPv4-in-IPv6 tunnel interface. The CE regards the BR as its DHCPv4-over-IPv6 server/relay, which is used to obtain its global IPv4 address allocation; its 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 one of the IPv6 addresses for DHCPv4 over IPv6 or negotiation of static configuration. The CE then uses the same IPv6 address 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 the 4over6 CE is simple. When receiving an IPv4-in-IPv6 packet, the CE just removes the IPv6 header and either hands it to a local upper layer or forwards it to the customer network according to the IPv4 destination address.

A CE runs a regular IPv4 Network Address and Port Translation (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 necessary when the CE is provisioned with an IPv4 prefix. In this case, detailed customer network planning is out of scope for this document.

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

6. 4over6 BR Behavior

The 4over6 BR maintains the bindings between the CE IPv6 address and CE IPv4 address (prefixes). The bindings are used to provide the correct encapsulation destination address for inbound IPv4 packets and also to 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 co-located 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 deletes it when the lease expires or a DHCP RELEASE message is received.
- o The BR can play the role of IPv6-Transport Relay Agent (TRA) as described in [DHCPv4-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 and uses the IPv4 destination address in the packet to look up the 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 the hairpinning of traffic between two CEs by performing decapsulation and re-encapsulation of packets.

In cases where the BR manages the global IPv4 address pool, the BR advertises the routing information of IPv4 addresses to the IPv4 Internet.

7. Fragmentation and Reassembly

The same considerations as those described in Sections 5.3 and 6.3 of [RFC6333] are taken into account for the CE and the BR, respectively.

8. DNS

The procedure described in Sections 5.5 and 6.4 of [RFC6333] is followed by the CE and the BR, respectively.

9. Security Considerations

The 4over6 BR implements methods to limit service only to registered customers. On the control plane, the BR allocates IPv4 addresses only to registered customers. The BR can filter on the IPv6 source addresses of incoming DHCP requests and only respond to the ones that are conveyed by registered IPv6 source addresses. But this doesn't work in situations where multi-homing is present. In the networks where Public 4over6 is deployed, multi-homing is disallowed to avoid this issue.

Alternatively, the BR can filter out the unregistered CE's requests during the DHCP process. For data packets, the BR does ingress filtering by looking up addresses in the IPv4-IPv6 address-binding table for the related matches as described in Section 6.

In the case of fallback to DS-Lite, security considerations in Section 11 of [RFC6333] are followed.

10. Contributors

The following are those who have made contributions 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

11. References

11.1. Normative References

- [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, December 1998.
- [RFC4925] Li, X., Dawkins, S., Ward, D., and A. Durand, "Softwire Problem Statement", RFC 4925, July 2007.
- [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.

11.2. Informative References

- [DHCPv4-IPv6]
Cui, Y., Wu, P., Wu, J., Lemon, T., and Q. Sun, "DHCPv4 over IPv6 Transport", Work in Progress, October 2013.
- [SOFTWARE-CPE]
Boucadair, M., Farrer, I., Perreault, S., Ed., and S. Sivakumar, Ed., "Unified IPv4-in-IPv6 Softwire CPE", Work in Progress, May 2013.
- [SOFTWARE-LW46]
Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture", Work in Progress, November 2013.

Authors' Addresses

Yong Cui
Tsinghua University
Beijing 100084
P.R.China
Phone: +86-10-6260-3059
EMail: yong@csnet1.cs.tsinghua.edu.cn

Jianping Wu
Tsinghua University
Beijing 100084
P.R.China
Phone: +86-10-6278-5983
EMail: jianping@cernet.edu.cn

Peng Wu
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