Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: January 5, 2012

T. Tsou
Huawei Technologies (USA)
C. Zhou
T. Taylor
Huawei Technologies
O. Troan
Cisco
Q. Chen
China Telecom

July 4, 2011

"Gateway-Initiated" 6rd
draft-tsou-gwinit-6rd-00

Abstract

This document proposes an alternative 6rd deployment model to that of RFC 5969. The basic 6rd model allows IPv6 hosts to gain access to IPv6 networks across an IPv4 access network using 6-in-4 tunnels. 6rd requires support by a device (the 6rd-CE) on the customer site, which must also be assigned an IPv4 address. The alternative model described in this document initiates the 6-in-4 tunnels from an operator-owned gateway collocated with the operator's IPv4 network edge, rather than from customer equipment. The advantages of this approach is that it requires no modification to customer equipment and avoids assignment of IPv4 addresses to customer equipment. The latter point means less pressure on IPv4 addresses in a high-growth environment.

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 January 5, 2012.

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

6rd ([RFC5969]) provides a transition tool for connecting IPv6 devices across an IPv4 network to an IPv6 network, at which point the packets can be routed natively. The network topology is shown in Figure 1.

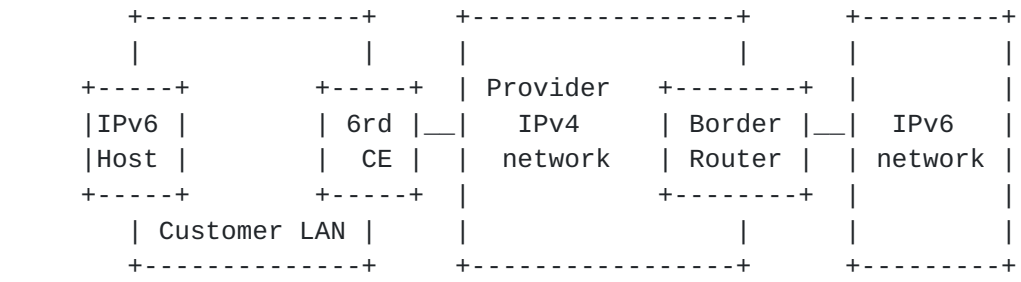

Figure 1: 6rd Deployment Topology

In Figure 1, the CE is the customer edge router. It is provisioned with a delegated IPv6 prefix, but also with an IPv4 address so that it is reachable through the IPv4 network. If public IPv4 address is provisioned to every customer, it will aggravate the pressure due to IPv4 address shortage for operators faced with a high rate of growth in the number of broadband subscribers to their network. It is out of scope of this document if private IPv4 address is provisioned.

2. Problem Statement

Consider an operator facing a high subscriber growth rate. As a result of this growth rate, the operator faces pressure on its stock of available public IPv4 addresses. For this reason, the operator is motivated to offer IPv6 access as quickly as possible.

The backbone network will be the first part of the operator's network to support IPv6. The metro network is not so easily upgraded to support IPv6 since many devices need to be modified and there may be some impact to existing services. Thus any means of providing IPv6 access has to minimize the changes required to devices in the metro network.

In contrast to the situation described for basic 6rd [RFC5569], the operator is assumed to be unable to manage IP devices on the customer premises. As a result, the operator cannot assume that any of these devices are capable of supporting 6rd.

If the customer equipment is in bridged mode and IPv6 is deployed to sites via a Service Provider's (SP's) IPv4 network, the IPv6-only

host needs a IPv6 address to visit the IPv6 service. In this scenario, 6to4 or 6RD can be used. However, each IPv6-only host may need one corresponding IPv4 address when using public IPv4 address in 6to4 or 6rd, which brings great address pressure to the operators.

If the customer equipment is in routing mode, the operator has an opportunity to avoid assigning IPv4 addresses to sites running IPv6 only. Some other means is available for routing IPv6 traffic through the IPv4 network to that site. The GateWay in the existing IPv4 access network should be updated to support IPv6. But the metro network does not need to be updated.

In 6rd scenario, reachability between CEs in 6RD should go to BR. But in this Gateway-initiated 6rd case, it does not need to go to BR which only needs gateway to gateway traffic. How the interaction between GW and GW works is for further elaboration.

3. Proposed Solution

For basic 6rd, the 6rd-CE described in [RFC5969] initiates the 6-in-4 tunnel to the Border Router to carry its IPv6 traffic. To avoid the requirement for customer premises equipment to fulfill this role, it is necessary to move the tunneling function to a network device. This document identifies a functional element termed the 6rd PE to perform this task. The functions of 6rd PE are:

- to generate and allocate gateway initiated 6rd delegated prefixes for IPv6-capable customer devices, as described in Section 3.1.

- to forward outgoing IPv6 packets through a tunnel to a Border Relay, which extracts and forwards them to an IPv6 network as for 6rd;

- to extract incoming IPv6 packets tunneled from the 6rd Border Relay and forward them to the correct user device.

In the proposed solution, there is only one tunnel initiated from each Gateway to the Border Router, which greatly reduces the number of tunnels the Border Router has to handle. The deployment scenario consistent with the problem statement in Section 2 collocates the Gateway with the IP edge of the access network. This is shown in Figure 2, and is the typical placement of the Broadband Network Gateway (BNG) in a fixed broadband network. By assumption, the metro network beyond the BNG is IPv4. Transport between the customer site and the Gateway is over layer 2.

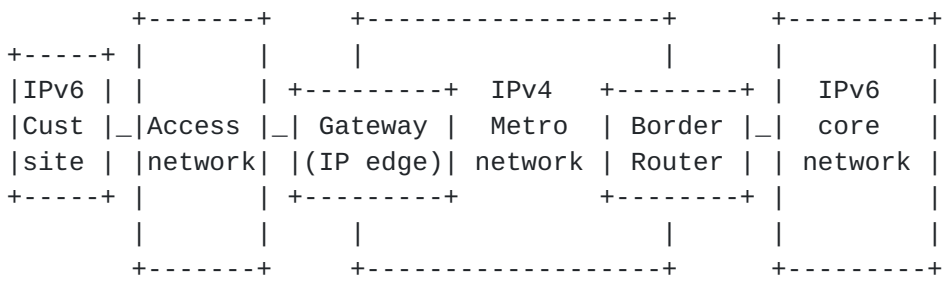

Figure 2: Gateway-Initiated 6rd At the IP Edge

The elements of the proposed solution are these:

- The IPv6 prefix assigned to the customer site contains the compressed IPv4 address of the network-facing side of the Gateway, plus a manually provisioned or Gateway-generated customer site identifier. This is illustrated in Figure 3 below.

- The Border Router is able to route incoming IPv6 packets to the correct Gateway by extracting the compressed Gateway address from the IPv6 destination address of the incoming packet, expanding it to a full 32-bit IPv4 address, and setting it as the destination address of the encapsulated packet.

- The Gateway can route incoming packets to the correct link after decapsulation using a mapping from either the full IPv6 prefix or the customer site identifier extracted from that prefix to the appropriate link.

3.1. Prefix Delegation

Referring back to Figure 2, prefix assignment to the customer equipment occurs in the normal fashion through the Gateway/IP edge, using either DHCPv6 or SLAAC. Figure 3 illustrates the structure of the assigned prefix, and how the components are derived, within the context of a complete address.

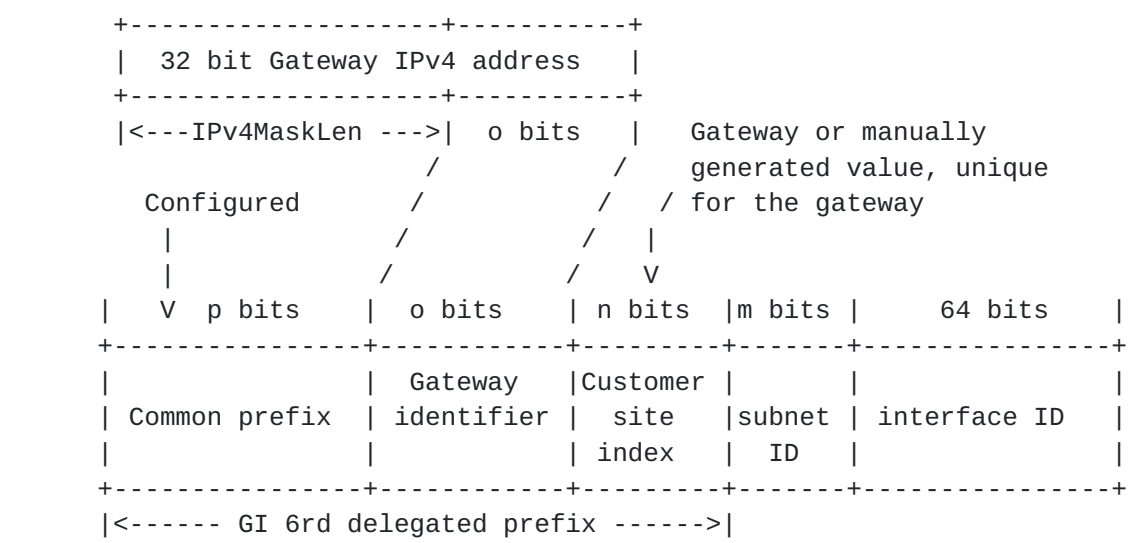

Figure 3: Gateway-Initiated 6rd Address Format for a Customer Site

The common prefix, i.e., the first p bits of the GI 6rd delegated prefix, is configured in the Gateway. This part of the prefix is common across multiple customers and multiple Gateways. Multiple common prefix values may be used in a network either for service separation or for scalability.

The Gateway Identifier is equal to the o low-order bits of the Gateway IPv4 address on the virtual link to the Border Router. The number of bits o is equal to 32 - IPv4MaskLen, where the latter is the length of the IPv4 prefix from which the Gateway IPv4 addresses are derived. The value of IPv4MaskLen is configured in both the Gateways and the Border Routers.

The Customer Site Index is effectively a sequence number assigned to an individual customer site served by the Gateway. The value of the index for a given customer site must be unique across the Gateway. The length n of the Customer Site Index is provisioned in the Gateway, and must be large enough to accommodate the number of customer sites that the Gateway is expected to serve.

To give a numerical example, consider a 6rd domain containing ten million IPv6-capable customer devices (a rather high number given that 6rd is meant for the early stages of IPv6 deployment). The estimated number of 6rd Gateways needed to serve this domain would be in the order of 3,300, each serving 30,000 customer devices. Assuming best-case compression for the Gateway addresses, the Gateway Identifier field has length o = 12 bits. If IPv6-in-IPv4 tunneling is being used, this best case is more likely to be achievable than it would be if the IPv4 addresses belonged to the customer devices. More controllably, the customer device index has length n = 15 bits.

Overall, these figures suggest that the length p of the common prefix can be 29 bits for a /56 delegated prefix, or 21 bits if /48 delegated prefixes need to be allocated.

3.2. Relevant Differences From Basic 6rd

A number of the points in [RFC5969] apply with the simple substitution of the Gateway for the 6rd CE. When it comes to configuration, the definition of IPv4MaskLen changes, and there are other differences as indicated in the previous section. Since special configuration of customer equipment is not required, the 6rd DHCPv6 option is inapplicable.

Since the link for the customer site to the network now extends only as far as the Gateway, Neighbour Unreachability Detection on the part of customer devices is similarly limited in scope.

3.3. Security Considerations

No change from [RFC5969].

3.4. IANA Considerations

This memo makes no request of IANA.

4. References

4.1. Normative References

[RFC5969] Townsley, W. and O. Troan, "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) -- Protocol Specification", RFC 5969, August 2010.

4.2. informative References

[RFC5569] Despres, R., "IPv6 Rapid Deployment on IPv4 Infrastructures (6rd)", RFC 5569, January 2010.

Authors' Addresses

Tina Tsou
Huawei Technologies (USA)
2330 Central Expressway
Santa Clara CA 95050
USA

Phone:
Email: tena@huawei.com

Cathy Zhou
Huawei Technologies
Bantian, Longgang District
Shenzhen 518129
P.R. China

Phone:
Email: cathyzhou@huawei.com

Tom Taylor
Huawei Technologies
1852 Lorraine Ave.t
Ottawa, Ontario K1H 6Z8
Canada

Phone:
Email: tom111.taylor@bell.net

Ole Troan
Cisco

Phone:
Email: ot@cisco.com

Qi Chen
China Telecom
109, Zhongshan Ave. West,
Tianhe District, Guangzhou 510630
P.R. China

Phone:
Email: chenqi.0819@gmail.com