Internet Engineering Task Force   T. Tsou
Internet-Draft   Huawei Technologies (USA)
Intended status: Informational   C. Zhou
Expires: September 6, 2012   T. Taylor
                                               Huawei Technologies
                                                            Q. Chen
                                                      China Telecom
                                                      March 5, 2012

                        "Gateway-Initiated" 6rd
                    draft-tsou-softwire-gwinit-6rd-06

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 are 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.

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.

```
    +--------------+    +----------------+     +---------+
    |              |    |                |     |         |
+-----+        +-----+  | Provider    +--------+ |         |
|IPv6 |        | 6rd |__|   IPv4      | Border |__| IPv6  |
|Host |        | CE  |  |  network    | Router |  | network |
+-----+        +-----+  |             +--------+ |         |
    | Customer LAN |    |                |     |         |
    +--------------+    +----------------+     +---------+
```
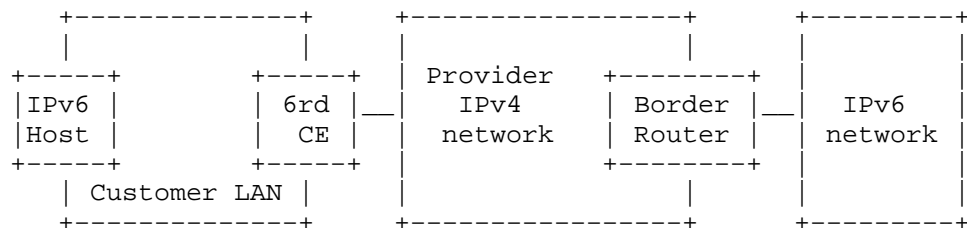
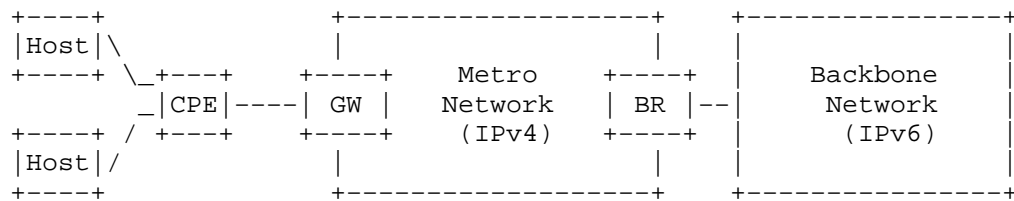   Figure 1: 6rd Deployment Topology

   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 a 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. The
   use of private addresses with 6rd avoids this particular difficulty,
   but brings other complications.

1.1. Requirements Language

   This document uses no requirements language.

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. Figure 2
   shows the sort of network situation envisioned in the present
   document.

```
+----+              +-----------------+    +---------------+
|Host|\             |                 |    |               |
+----+ \_+---+    +----+   Metro    +----+ |   Backbone    |
       _|CPE|----| GW |  Network    | BR |--|   Network     |
+----+ / +---+    +----+  (IPv4)     +----+ |    (IPv6)     |
|Host|/             |                 |    |               |
+----+              +-----------------+    +---------------+

Host = IPv6 customer host device
CPE  = customer edge device (customer-provided)
GW   = provider edge device (Gateway)
BR   = border router (dual stack)
```

Specialized GW and BR functions are described in the next section.

Figure 2: Typical Network Scenario For IPv6 Transition

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 have no control over the capabilities of the
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 puts great address pressure on the operators.

If the CPE in the above figure is acting in bridging mode, each host
behind it needs to be directly assigned an IPv6 prefix so it can
access IPv6 services. If the CPE is acting in routing mode, only the
CPE needs to be assigned an IPv6 prefix, and it delegates prefixes to
the hosts behind it.

If the Gateway supports IPv4 only, then an IPv4 address must also be
assigned to each host (bridging mode) or to the CPE (routing mode).
Both cases, but bridging mode in particular, put pressure on the
provider's stock of IPv4 addresses.

If the Gateway is dual stack, an arrangement may be possible whereby
all communication between the Gateway and the customer site uses IPv6
and the need to assign IPv4 addresses to customer devices is avoided.
A possible solution is presented in the next section.

3. Proposed Solution

For basic 6rd [RFC5969], the 6rd CE initiates the 6-in-4 tunnel to
the 6rd Border Relay 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 Gateway
to perform this task. In what follows, the 6rd Gateway and 6rd
Border Relay are referred to simply as the Gateway and Border Relay
respectively.

The functions of Gateway are:

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

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

o   to extract incoming IPv6 packets tunneled from the 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 Relay, which greatly reduces the number of
tunnels the Border Relay 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 above, 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.

The elements of the proposed solution are these:

o   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.

o   The Border Relay 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.

o 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.

```
      +--------------------+-----------+
      |  32 bit Gateway IPv4 address   |
      +--------------------+-----------+
      |<---IPv4MaskLen --->|  o bits   |     Gateway or manually
                          /         /       generated value, unique
         Configured      /         /     / for the gateway
         |              /         /    |
         |             /         /     V
      |  V  p bits    |  o bits  | n bits  |m bits |    64 bits    |
      +---------------+----------+---------+-------+---------------+
      |               | Gateway  |Customer |       |               |
      | Common prefix | identifier|  site  |subnet | interface ID  |
      |               |          | index   |  ID   |               |
      +---------------+----------+---------+-------+---------------+
      |<------ GI 6rd delegated prefix ------>|
```
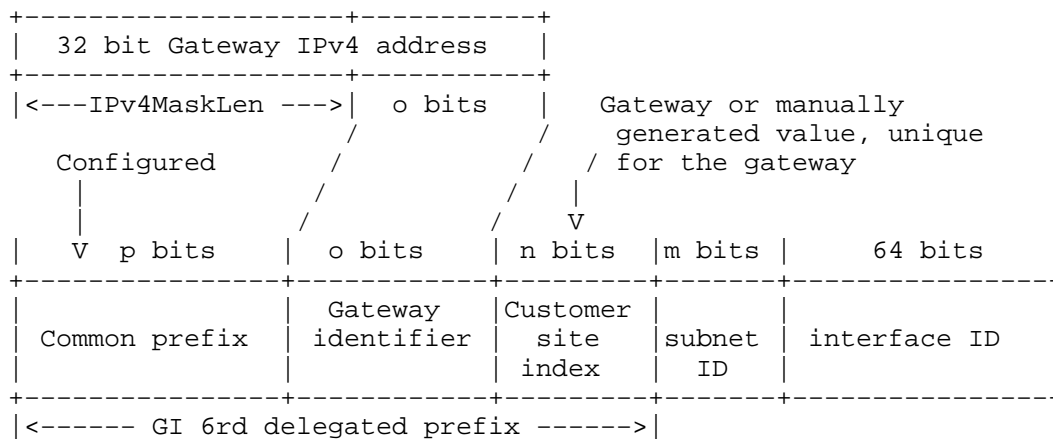
Figure 3: Gateway-Initiated 6rd Address Format for a Customer Site

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 Relay. 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 Relays.

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.

4. IANA Considerations

This memo includes no request to IANA.

5. Security Considerations

No change from [RFC5969].

6. Acknowledgements

   Thanks to Ole Troan for his technical comments on an early version of
   this document.

7. References

7.1. Normative References

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

7.2. 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: Tina.Tsou.Zouting@huawei.com

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

   Phone:
   Email: cathy.zhou@huawei.com

Tom Taylor
Huawei Technologies
Ottawa, Ontario
Canada

Phone:
Email: tom.taylor.stds@gmail.com

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

Phone:
Email: chenqi.0819@gmail.com