Network Working Group   S. Jiang
Internet Draft   D. Guo
Intended status: Informational   Huawei Technologies Co., Ltd
Expires: July 4, 2011   B. Carpenter
                                          University of Auckland
                                                January 4, 2011

         An Incremental Carrier-Grade NAT (CGN) for IPv6 Transition
               draft-ietf-v6ops-incremental-cgn-03.txt

Status of this Memo

Copyright Notice

Abstract

   Global IPv6 deployment was slower than originally expected. As IPv4
   address exhaustion approaches, IPv4 to IPv6 transition issues become
   more critical and less tractable. Host-based transition mechanisms

used in dual stack environments cannot meet all transition
requirements. Most end users are not sufficiently expert to configure
or maintain host-based transition mechanisms. Carrier-Grade NAT (CGN)
devices with integrated transition mechanisms can reduce the
operational changes required during the IPv4 to IPv6 migration or
coexistence period.

This document proposes an incremental CGN approach for IPv6
transition. It can provide IPv6 access services for IPv6 hosts and
IPv4 access services for IPv4 hosts, while leaving much of a legacy
ISP network unchanged during the initial stage of IPv4 to IPv6
migration. Unlike CGN alone, incremental CGN also supports and
encourages smooth transition towards dual-stack or IPv6-only ISP
networks. An integrated configurable CGN device and an adaptive Home
Gateway (HG) device are described. Both are re-usable during
different transition phases, avoiding multiple upgrades. This enables
IPv6 migration to be incrementally achieved according to real user
requirements.

Table of Contents

1. Introduction

   Global IPv6 deployment did not happen as was forecast 10 years ago.
   Network providers were hesitant to make the first move while IPv4 was
   and is still working well. However, IPv4 address exhaustion is
   imminent. The dynamically-updated IPv4 Address Report [IPUSAGE] has
   analyzed this issue. It predicts early 2011 for IANA unallocated
   address pool exhaustion and middle 2012 for RIR unallocated address
   pool exhaustion. Based on this fact, the Internet industry appears to
   have reached consensus that global IPv6 deployment is inevitable and
   has to be done expeditiously.

   IPv4 to IPv6 transition issues therefore become more critical and
   complicated for the approaching global IPv6 deployment. Host-based
   transition mechanisms alone are not able to meet the requirements in
   all cases. Therefore, network-based supporting functions and/or new
   transition mechanisms with simple user-side operation are needed.

   Carrier-Grade NAT (CGN) [I-D.nishitani-cgn], also called NAT444 CGN
   or Large Scale NAT, compounds IPv4 operational problems when used
   alone, but does nothing to encourage IPv4 to IPv6 transition.
   Deployment of NAT444 CGN allows ISPs to delay the transition, and
   therefore causes double transition costs (once to add CGN, and again
   to support IPv6).

   CGN deployments that integrate multiple transition mechanisms can
   simplify the operation of end user services during the IPv4 to IPv6
   migration and coexistence periods. CGNs are deployed on the network
   side and managed/maintained by professionals. On the user side, new
   Home Gateway (HG) devices may be needed too. They may be provided by
   network providers, depending on the specific business model. Dual-
   stack lite [I-D.ietf-softwire-dual-stack-lite], also called DS-Lite,
   is a CGN-based solution that supports transition, but it requires the
   ISP to upgrade its network to IPv6 immediately. Many ISPs hesitate to
   do this as the first step. Theoretically, DS-Lite can be used with
   double encapsulation (IPv4-in-IPv6-in-IPv4) but this seems even less
   likely to be accepted by an ISP and is not discussed in this
   document.

   This document proposes an incremental CGN approach for IPv6
   transition. It does not define any new protocols or mechanisms, but
   describes how to combine existing proposals in an incremental
   deployment. The approach is similar to DS-Lite, but the other way
   around. It mainly combines v4-v4 NAT with v6-over-v4 tunneling
   functions. It can provide IPv6 access services for IPv6-enabled hosts
   and IPv4 access services for IPv4 hosts, while leaving most of legacy

IPv4 ISP networks unchanged. The deployment of this technology does
not affect legacy IPv4 hosts with global IPv4 addresses at all. It is
suitable for the initial stage of IPv4 to IPv6 migration. It also
supports transition towards dual-stack or IPv6-only ISP networks.

A smooth transition mechanism is also described in this document. It
introduces an integrated configurable CGN device and an adaptive HG
device. Both CGN and HG are re-usable devices during different
transition periods, so they do not need to be replaced as the
transition proceeds. This enables IPv6 migration to be incrementally
achieved according to the real user requirements.

2. An Incremental CGN Approach

Today, most consumers primarily use IPv4. Network providers are
starting to provide IPv6 access services for end users. At the
initial stage of IPv4 to IPv6 migration, IPv4 connectivity and
traffic would continue to represent the majority of traffic for most
ISP networks. ISPs would like to minimize the changes to their IPv4
networks. Switching the whole ISP network into IPv6-only would be
considered as a radical strategy. Switching the whole ISP network to
dual stack is less radical, but introduces operational costs and
complications. Although some ISPs have successfully deployed dual
stack networks, others prefer not to do this as their first step in
IPv6. However, they currently face two urgent pressures - to
compensate for an immediate shortage of IPv4 addresses by deploying
some method of address sharing, and to prepare actively for the use
of deployment of IPv6 address space and services. ISPs facing only
one pressure out of two could adopt either CGN (for shortage of IPv4
addresses) or 6rd (to provide IPv6 connectivity services). The
approach described in this draft is intended to address both of these
pressures at the same time by combining v4-v4 CGN with v6-over-v4
tunneling technologies.

2.1. Incremental CGN Approach Overview

The incremental CGN approach we propose is illustrated as the
following figure.

```
                        +-------------+
                        |IPv6 Internet|
                        +-------------+
                               |
                    +---------------+----------+
  +-----+    +--+    |   IPv4 ISP  +--+--+     |     +--------+
  |v4/v6|---|DS|=======+===========| CGN |-------+---|  IPv4  |
  |Host |   |HG|    |   Network  +-----+   |     |   |Internet|
  +-----+    +--+    +-------------------+---+     +--------+
               _ _ _ _ _ _ _ _ _ _            |
        ()_6_over_4_ _t_u_n_n_e_l_()  +--------------------+
                                      | Existing IPv4 hosts |
                                      +--------------------+
```
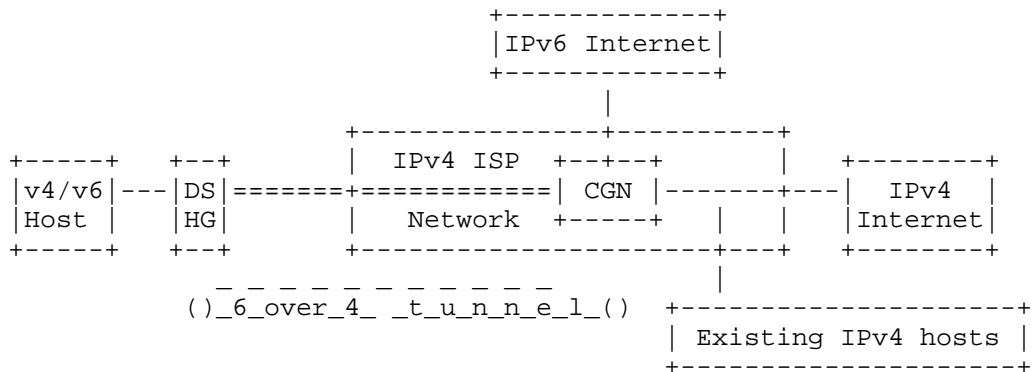
   Figure 1: incremental CGN approach with IPv4 ISP network

   DS HG = Dual-Stack Home Gateway (CPE - Customer Premises Equipment).

   As shown in the above figure, the ISP has not significantly changed
   its IPv4 network. This approach enables IPv4 hosts to access the IPv4
   Internet and IPv6 hosts to access the IPv6 Internet. A dual stack
   host is treated as an IPv4 host when it uses IPv4 access service and
   as an IPv6 host when it uses an IPv6 access service. In order to
   enable IPv4 hosts to access the IPv6 Internet and IPv6 hosts to
   access IPv4 Internet, NAT64 can be integrated with the CGN; see
   Section 2.6 for details regarding IPv4/IPv6 intercommunication. The
   integration of such mechanisms is out of scope for this document.

   Two types of device need to be deployed in this approach: a dual-
   stack home gateway (HG), and a dual-stack CGN. The dual-stack home
   gateway integrates both IPv6 and IPv4 forwarding and v6-over-v4
   tunneling functions. It should follow the requirements of
   [I-D.ietf-v6ops-ipv6-cpe-router], including IPv6 prefix delegation.
   It may integrate v4-v4 NAT functionality, too. The dual-stack CGN
   integrates v6-over-v4 tunneling and v4-v4 CGN functions, as well as
   standard IPv6 and IPv4 routing

   The approach does not require any new mechanisms for IP packet
   forwarding or encapsulation or decapsulation at tunnel end points.
   The following sections describe how the HG and the incremental CGN
   interact.

2.2. Choice of tunneling technology

   In principle, this model will work with any form of tunnel between
   the dual-stack HG and the dual-stack CGN. However, tunnels that
   require individual configuration are clearly undesirable because of

their operational cost. Configured tunnels based directly on
[RFC4213] are therefore not suitable. A tunnel broker according to
[RFC3053] would also have high operational costs and be unsuitable
for home users.

6rd [RFC5569, RFC5969] technology appears suitable to support v6-
over-v4 tunneling with low operational cost. GRE [RFC2784] with an
additional auto-configuration mechanism is also able to support v6-
over-v4 tunneling. Other tunneling mechanisms such as 6over4
[RFC2529], 6to4 [RFC3056], the Intra-Site Automatic Tunnel Addressing
Protocol (ISATAP) [RFC5214] or Virtual Enterprise Traversal (VET)
[RFC5558] could be considered. If the ISP has an entirely MPLS
infrastructure between the HG and the dual-stack CGN, it would also
be possible to use a 6PE [RFC4798] tunnel directly over MPLS. This
would, however, only be suitable for an advanced HG that is unlikely
to be found as a consumer device, and is not further discussed here.
To simplify the discussion, we assume the use of 6rd.

2.3. Behavior of Dual-stack Home Gateway

When a dual-stack home gateway receives a data packet from a host, it
will determine whether the packet is an IPv4 or IPv6 packet. The
packet will be handled by an IPv4 or IPv6 stack as appropriate. For
IPv4, and in the absence of v4-v4 NAT on the HG, the stack will
simply forward the packet to the CGN, which will normally be the IPv4
default router. If v4-v4 NAT is enabled, the HG translates the packet
source address from a HG-scope private IPv4 address into a CGN-scope
IPv4 address, performs port mapping if needed, and then forwards the
packet towards the CGN. The HG records the v4-v4 address and port
mapping information for inbound packets, like any other NAT.

For IPv6, the HG needs to encapsulate the data into an IPv4 tunnel
packet, which has the dual-stack CGN as the IPv4 destination. The HG
sends the new IPv4 packet towards the CGN, for example using 6rd.

If the HG is linked to more than one CGN, it will record the mapping
information between the tunnel and the source IPv6 address for
inbound packets. Detailed considerations for the use of multiple CGNs
by one HG are for further study.

IPv4 packets from the CGN, and encapsulated IPv6 packets from the
CGN, will be translated or decapsulated according to the stored
mapping information and forwarded to the customer side of the HG.

2.4. Behavior of Dual-stack CGN

   When a dual-stack CGN receives an IPv4 data packet from a dual-stack
   home gateway, it will determine whether the packet is a normal IPv4
   packet, which is non-encapsulated, or a v6-over-v4 tunnel packet
   addressed to a tunnel end point within the CGN. For a normal IPv4
   packet, the CGN translates the packet source address from a CGN-scope
   IPv4 address into a public IPv4 address, performs port mapping if
   necessary, and then forwards it normally to the IPv4 Internet. The
   CGN records the v4-v4 address and port mapping information for
   inbound packets, just like a normal NAT does. For a v6-over-v4 tunnel
   packet, the tunnel end point within the CGN will decapsulate it into
   the original IPv6 packet and then forward it to the IPv6 Internet.
   The CGN records the mapping information between the tunnel and the
   source IPv6 address for inbound packets.

   Depending on the deployed location of the CGN, it may use a further
   v6-over-v4 tunnel to connect to the IPv6 Internet.

   Packets from the IPv4 Internet will be appropriately translated by
   the CGN and forwarded to the HG, and packets from the IPv6 Internet
   will be tunneled to the appropriate HG, using the stored mapping
   information as necessary.

2.5. Impact for existing hosts and unchanged networks

   This approach does not affect the unchanged parts of ISP networks at
   all. Legacy IPv4 ISP networks and their IPv4 devices remain in use.
   The existing IPv4 hosts, shown as the lower right box in Figure 1,
   either having global IPv4 addresses or behind v4-v4 NAT, can connect
   to the IPv4 Internet as it is now. These hosts, if they are upgraded
   to become dual-stack hosts, can access the IPv6 Internet through the
   IPv4 ISP network by using IPv6-over-IPv4 tunnel technologies. (See
   section 2.7 for a comment on MTU size.)

2.6. IPv4/IPv6 intercommunication

   IPv6-only public services are not expected as long as there is
   significant IPv4-only customer base in the world, for obvious
   commercial reasons. However, IPv4/IPv6 intercommunication may become
   issues in many scenarios.

   The IETF is expected to standardize a recommended IPv4/IPv6
   translation algorithm, sometimes referred to as NAT64. It is
   specified in

o "Framework for IPv4/IPv6 Translation"
   [I-D.ietf-behave-v6v4-framework]
o "IPv6 Addressing of IPv4/IPv6 Translators" [RFC6052]
o "DNS64: DNS extensions for Network Address Translation from IPv6
   Clients to IPv4 Servers" [I-D.ietf-behave-dns64]
o "IP/ICMP Translation Algorithm" [I-D.ietf-behave-v6v4-xlate]
o "Stateful NAT64: Network Address and Protocol Translation from
   IPv6 Clients to IPv4 Servers"
   [I-D.ietf-behave-v6v4-xlate-stateful]
o "An FTP ALG for IPv6-to-IPv4 translation" [I-D.ietf-behave-ftp64]

The service, as described in the IETF's "Guidelines for Using IPv6
Transition Mechanisms during IPv6 Deployment"
[I-D.arkko-ipv6-transition-guidelines], provides for stateless
translation between hosts in an IPv4-only domain or which offer an
IPv4-only service and hosts with an IPv4-embedded IPv6 address in an
IPv6-only domain. It additionally provide access from IPv6 hosts with
general format addresses to hosts in an IPv4-only domain or which
offer an IPv4-only service. It does not provide any-to-any
translation. One result is that client-only hosts in the IPv6 domain
gain access to IPv4 services through stateful translation. Another
result is that the IPv6 network operator has the option of moving
servers into the IPv6-only domain while retaining accessibility for
IPv4-only clients, through stateless translation of an IPv4-embedded
IPv6 address.

Also, "Architectural Implications of NAT" [RFC2993] applies across
the service just as in IPv4/IPv4 translation: apart from the fact
that a system with an IPv4-embedded IPv6 address is reachable across
the NAT, which is unlike IPv4, any assumption on the application's
part that a local address is meaningful to a remote peer, and any use
of an IP address literal in the application payload, is a source of
service issues. In general, the recommended mitigation for this is

o Ideally, applications should use DNS names rather than IP address
   literals in URLs, URIs, and referrals, and in general be network
   layer agnostic.
o If they do not, the network may provide a relay or proxy that
   straddles the domains. For example, an SMTP MTA with both IPv4
   and IPv6 connectivity handles IPv4/IPv6 translation cleanly at the
   application layer.

2.7. Discussion

For IPv4 traffic, the incremental CGN approach inherits all the
problems of CGN address sharing techniques
[I-D.ietf-intarea-shared-addressing-issues] (e.g., scaling, and the

difficulty of supporting well-known ports for inbound traffic).
Application layer problems created by double NAT are beyond the scope
of this document.

For IPv6 traffic, a user behind the DS HG will see normal IPv6
service. We observe that an IPv6 tunnel MTU of at least 1500 bytes
would ensure that the mechanism does not cause excessive
fragmentation of IPv6 traffic nor excessive IPv6 path MTU discovery
interactions. This, and the absence of NAT problems for IPv6, will
create an incentive for users and application service providers to
prefer IPv6.

ICMP filtering [RFC4890] may be included as part of CGN functions.

3. Smooth transition towards IPv6 infrastructure

Transition from pure NAT444 CGN or 6rd to the incremental CGN
approach is straightforward. The HG and CGN devices and their
locations do not have to be changed; only software upgrading may be
needed. In the ideal model, described below, even software upgrading
is not necessary; reconfiguration of the devices is enough. NAT444
CGN solves the public address shortage issues in the current IPv4
infrastructure. However, it does not contribute towards IPv6
deployment at all. The incremental CGN approach can inherit NAT444
CGN functions while providing overlay IPv6 services. 6rd mechanisms
can also transform smoothly into this incremental CGN model. However,
the home gateways need to be upgraded correspondingly to perform the
steps described below

The incremental CGN can also easily be transitioned to an IPv6-
enabled infrastructure, in which the ISP network is either dual-stack
or IPv6-only.

If the ISP prefers to move to dual-stack routing, the HG should
simply switch off its v6-over-v4 function when it observes native
IPv6 RA or DHCPv6 traffic, and then forward both IPv6 and IPv4
traffic directly, while the dual-stack CGN keeps only its v4-v4 NAT
function.

However, we expect ISPs to choose the approach described as
incremental CGN here because they intend to avoid dual-stack routing,
and to move incrementally from IPv4-only routing to IPv6-only
routing. In this case, the ideal model for the incremental CGN
approach is that of an integrated configurable CGN device and an
adaptive HG device. The integrated CGN device will be able to support
multiple functions, including NAT444 CGN, 6rd router (or an
alternative tunneling mechanism), DS-Lite, and dual-stack forwarding.

The HG has to integrate the corresponding functions, and be able to detect relevant incremental changes on the CGN side. Typically the HG will occasionally poll the CGN to discover which features are operational. For example, starting from a pure IPv4-only scenario (in which the HG treats the CGN only as an IPv4 default router), the HG would discover by infrequent polling when 6rd became available. The home user would then acquire an IPv6 prefix. At a later stage, the HG would observe the appearance of native IPv6 Route Advertisement messages or DHCPv6 messages to discover the availability of an IPv6 service including DS-Lite. Thus, when an ISP decides to switch from incremental CGN to DS-Lite CGN, only a configuration change or a minor software update is needed on the CGNs. The home gateway would detect this change and switch automatically to DS-Lite mode. The only impact on the home user will be to receive a different IPv6 prefix.

In the smooth transition model, both CGN and HG are re-usable devices during different transition periods. This avoids potential multiple upgrades. Different regions of the same ISP network may be at different stages of the incremental process, using identical equipment but with different configurations of the incremental CGN devices in each region. Thus, IPv6 migration may be incrementally achieved according to the real ISP and customer requirements.

4. Security Considerations

Security issues associated with NAT have been documented in [RFC2663] and [RFC2993]. Security issues for large-scale address sharing, including CGN, are discussed in [I-D.ietf-intarea-shared-addressing-issues]. The present specification does not introduce any new features to CGN itself, and hence no new security considerations. Security issues for 6rd are documented in [RFC5569, RFC5969] and those for DS-Lite in [I-D.ietf-softwire-dual-stack-lite].

Since the tunnels proposed here exist entirely within a single ISP network, between the HG/CPE and the CGN, the threat model is relatively simple. [RFC4891] describes how to protect tunnels using IPsec, but an ISP could reasonably deem its infrastructure to provide adequate security without the additional protection and overhead of IPsec. The intrinsic risks of tunnels are described in [I-D.ietf-v6ops-tunnel-security-concerns], which recommends that tunneled traffic should not cross border routers. The incremental CGN approach respects this recommendation. To avoid other risks caused by tunnels, it is important that any security mechanisms based on packet inspection, and any implementation of ingress filtering, are applied to IPv6 packets after they have been decapsulated by the CGN. The dual-stack home gateway will need to provide basic security

functionality for IPv6 [I-D.ietf-v6ops-cpe-simple-security]. Other aspects are described in [RFC4864].

5. IANA Considerations

   This draft does not request any IANA action.

6. Acknowledgements

   Useful comments were made by Fred Baker, Dan Wing, Fred Templin, Seiichi Kawamura, Remi Despres, Janos Mohacsi, Mohamed Boucadair, Shin Miyakawa, Joel Jaeggli, Jari Arkko, Tim Polk, Sean Turner and other members of the IETF V6OPS working group.

7. Change Log [RFC Editor please remove]

   draft-jiang-incremental-cgn-00, original version, 2009-02-27

   draft-jiang-v6ops-incremental-cgn-00, revised after comments at IETF74, 2009-04-23

   draft-jiang-v6ops-incremental-cgn-01, revised after comments at v6ops mailing list, 2009-06-30

   draft-jiang-v6ops-incremental-cgn-02, remove normative parts (to be documented in other WGs), 2009-07-06

   draft-jiang-v6ops-incremental-cgn-03, revised after comments at v6ops mailing list, 2009-09-24

   draft-ietf-v6ops-incremental-cgn-00, accepted as v6ops wg document, 2009-11-17

   draft-ietf-v6ops-incremental-cgn-01, revised after comments at v6ops mailing list, 2010-06-21

   draft-ietf-v6ops-incremental-cgn-02, revised after comments at v6ops WGLC, 2010-10-11

   draft-ietf-v6ops-incremental-cgn-03, revised according comments from IESG, 2011-1-4

8. References

8.1. Normative References

    [RFC2529] B. Carpenter, and C. Jung, "Transmission of IPv6 over IPv4
              Domains without Explicit Tunnels", RFC2529, March 1999.

    [RFC2784] D. Farinacci, T. Li, S. Hanks, D. Meyer and P. Traina,
              "Generic Routing Encapsulation (GRE)", RFC 2784, March
              2000.

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

    [RFC5969] W. Townsley and O. Troan, "IPv6 via IPv4 Service Provider
              Networks '6rd'", RFC5969, May 2010.

8.2. Informative References

    [RFC2663] P. Srisuresh and M. Holdrege, "IP Network Address
              Translator (NAT) Terminology and Considerations", RFC 2663,
              August 1999.

    [RFC2993] T. Hain, "Architectural Implications of NAT", RFC 2993,
              November 2000.

    [RFC3053] A. Durand, P. Fasano, I. Guardini and D. Lento, "IPv6
              Tunnel Broker", RFC 3053, January 2001.

    [RFC3056] B. Carpenter and K. Moore, "Connection of IPv6 Domains via
              IPv4 Clouds", RFC 3056, February 2001.

    [RFC4213] E. Nordmark and R. Gilligan, "Basic Transition Mechanisms
              for IPv6 Hosts and Routers", RFC 4213, October 2005.

    [RFC4798] J. De Clercq, D. Ooms, S. Prevost and F. Le Faucheur,
              "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider
              Edge Routers (6PE)", RFC 4798, February 2007.

    [RFC4864] G. Van de Velde, T. Hain, R. Droms, B. Carpenter and E.
              Klein, "Local Network Protection for IPv6", RFC 4864, May
              2007.

    [RFC4890] E. Davies and J. Mohacsi, "Recommendations for Filtering
              ICMPv6 Messages in Firewalls", RFC 4890, May 2007.

   [RFC4891] R. Graveman, "Using IPsec to Secure IPv6-in-IPv4 Tunnels",
            RFC4891, May 2007.

   [RFC5214] F. Templin, T. Gleeson and D. Thaler, "Intra-Site Automatic
            Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.

   [RFC5558] F. Templin, "Virtual Enterprise Traversal (VET)", RFC 5558,
            February 2010.

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

   [IPUSAGE] G. Huston, IPv4 Address Report, March 2009,
            http://www.potaroo.net/tools/ipv4/index.html.

   [I-D.ietf-softwire-dual-stack-lite]
            A. Durand, "Dual-stack lite broadband deployments post IPv4
            exhaustion", draft-ietf-softwire-dual-stack-lite, work in
            progress.

   [I-D.ietf-v6ops-ipv6-cpe-router]
            H. Singh, W. Beebee, C. Donley, B. Stark and O. Troan,
            "IPv6 CPE Router Recommendations", draft-ietf-v6ops-ipv6-
            cpe-router, work in progress.

   [I-D.ietf-v6ops-cpe-simple-security]
            J. Woodyatt, "Recommended Simple Security Capabilities in
            Customer Premises Equipment for Providing Residential IPv6
            Internet Service", draft-ietf-v6ops-cpe-simple-security,
            work in progress.

   [I-D.ietf-behave-v6v4-xlate-stateful]
            M. Bagnulo, P. Matthews and I. van Beijnum, "NAT64: Network
            Address and Protocol Translation from IPv6 Clients to IPv4
            Servers", draft-ietf-behave-v6v4-xlate-stateful, work in
            progress.

   [I-D.ietf-intarea-shared-addressing-issues]
            M. Ford, et al, "Issues with IP Address Sharing", draft-
            ietf-intarea-shared-addressing-issues, work in progress.

   [I-D.nishitani-cgn]
            I. Yamagata, T. Nishitani, S. Miyahawa, A. nakagawa and H.
            Ashida, "Common requirements for IP address sharing
            schemes", draft-nishitani-cgn, work in progress.

[I-D.arkko-ipv6-transition-guidelines]
          Arkko, J. and F. Baker, "Guidelines for Using IPv6
          Transition Mechanisms during IPv6 Deployment", draft-arkko-
          ipv6-transition-guidelines, work in progress.

[I-D.ietf-behave-dns64]
          Bagnulo, M., Sullivan, A., Matthews, P., and I. Beijnum,
          "DNS64: DNS extensions for Network Address Translation from
          IPv6 Clients to IPv4 Servers", draft-ietf-behave-dns64,
          work in progress.

[I-D.ietf-behave-ftp64]
          Beijnum, I., "An FTP ALG for IPv6-to-IPv4 translation",
          draft-ietf-behave-ftp64, work in progress.

[I-D.ietf-behave-v6v4-framework]
          Baker, F., Li, X., Bao, C., and K. Yin, "Framework for
          IPv4/IPv6 Translation", draft-ietf-behave-v6v4-framework,
          work in progress.

[I-D.ietf-behave-v6v4-xlate]
          Li, X., Bao, C., and F. Baker, "IP/ICMP Translation
          Algorithm", draft-ietf-behave-v6v4-xlate, work in
          progress.

Author's Addresses

    Sheng Jiang
    Huawei Technologies Co., Ltd
    Huawei Building, No.3 Xinxi Rd.,
    Shang-Di Information Industry Base, Hai-Dian District, Beijing 100085
    P.R. China
    Email: shengjiang@huawei.com

    Dayong Guo
    Huawei Technologies Co., Ltd
    Huawei Building, No.3 Xinxi Rd.,
    Shang-Di Information Industry Base, Hai-Dian District, Beijing 100085
    P.R. China
    Email: guoseu@huawei.com

    Brian Carpenter
    Department of Computer Science
    University of Auckland
    PB 92019
    Auckland, 1142
    New Zealand
    Email: brian.e.carpenter@gmail.com