Basic Requirements for IPv6 Customer Edge Routers
draft-ietf-v6ops-ipv6-cpe-router-07

Abstract

   This document specifies requirements for an IPv6 Customer Edge (CE)
   router. Specifically, the current version of this document focuses
   on the basic provisioning of an IPv6 CE router and the provisioning
   of IPv6 hosts attached to it.

Table of Contents

1. Introduction

   This document defines basic IPv6 features for a residential or small
   office router referred to as an IPv6 CE router. Typically these
   routers also support IPv4.

   Mixed environments of dual-stack hosts and IPv6-only hosts (behind
   the CE router) can be more complex if the IPv6-only devices are using
   a translator to access IPv4 servers [I-D.ietf-behave-v6v4-framework].
   Support for such mixed environments is not in scope of this document.

   This document specifies how an IPv6 CE router automatically
   provisions its WAN interface, acquires address space for provisioning
   of its LAN interfaces and fetches other configuration information
   from the service provider network. Automatic provisioning of more
   complex topology than a single router with multiple LAN interfaces is
   out of scope for this document.

   See [RFC4779] for a discussion of options available for deploying
   IPv6 in Service Provider access networks.

1.1. Requirements Language

   The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
   "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
   document are to be interpreted as described in RFC 2119 [RFC2119].

2. Terminology

   End-user Network   one or more links attached to the IPv6 CE
                             router that connect IPv6 hosts.

   IPv6 Customer Edge router a node intended for home or small office
                             use which forwards IPv6 packets not
                             explicitly addressed to itself. The IPv6
                             CE router connects the end-user network to
                             a service provider network.

   IPv6 host   any device implementing an IPv6 stack
                             receiving IPv6 connectivity through the
                             IPv6 CE router

   LAN interface   an IPv6 CE router's attachment to a link in
                             the end-user network. Examples are
                             Ethernets (simple or bridged), 802.11
                             wireless or other LAN technologies. An
                             IPv6 CE router may have one or more network

layer LAN Interfaces.

Service Provider   an entity that provides access to the
                          Internet. In this document, a Service
                          Provider specifically offers Internet
                          access using IPv6, and may also offer IPv4
                          Internet access. The Service Provider can
                          provide such access over a variety of
                          different transport methods such as DSL,
                          cable, wireless, and others.

WAN interface   an IPv6 CE router's attachment to a link
                          used to provide connectivity to the Service
                          Provider network; example link technologies
                          include Ethernets (simple or bridged), PPP
                          links, Frame Relay, or ATM networks as well
                          as Internet-layer (or higher-layer)
                          "tunnels", such as tunnels over IPv4 or
                          IPv6 itself.

3. Architecture

3.1. Current IPv4 End-user Network Architecture

   An end-user network will likely support both IPv4 and IPv6. It is
   not expected that an end-user will change their existing network
   topology with the introduction of IPv6. There are some differences
   in how IPv6 works and is provisioned which has implications for the
   network architecture. A typical IPv4 end-user network consist of a
   "plug and play" router with NAT functionality and a single link
   behind it, connected to the Service Provider network.

   A typical IPv4 NAT deployment by default blocks all incoming
   connections. Opening of ports is typically allowed using UPnP IGD
   [UPnP-IGD] or some other firewall control protocol.

   Another consequence of using private address space in the end-user
   network is that it provides stable addressing, i.e. it never changes
   even when you change Service Providers, and the addresses are always
   there even when the WAN interface is down or the customer edge router
   has not yet been provisioned.

   Rewriting addresses on the edge of the network also allows for some
   rudimentary multi-homing; even though using NATs for multi-homing
   does not preserve connections during a fail-over event [RFC4864].

   Many existing routers support dynamic routing, and advanced end users

can build arbitrary, complex networks using manual configuration of
address prefixes combined with a dynamic routing protocol.

3.2. IPv6 End-user Network Architecture

The end-user network architecture for IPv6 should provide equivalent
or better capabilities and functionality than the current IPv4
architecture.

The end-user network is a stub network. Figure 1 illustrates the
model topology for the end-user network.

```
            An example of a typical end-user network.


                +-------+-------+                    \
                |    Service    |                     \
                |    Provider   |                      | Service
                |     Router    |                      | Provider
                +-------+-------+                      | network
                        |                             /
                        | Customer                   /
                        | Internet connection       /
                        |
                 +------+--------+                   \
                 |      IPv6     |                    \
                 | Customer Edge |                     \
                 |     Router    |                     /
                 +---+-------+-+-+                    /
          Network A      |     |  Network B    | End-User
        ---+------------+----+-   --+--+------------+---  |network(s)
           |            |           |            |     \
      +----+-----+ +-----+----+   +----+-----+ +-----+----+  \
      |IPv6 Host | |IPv6 Host |   | IPv6 Host| |IPv6 Host |  /
      |          | |          |   |          | |          | /
      +----------+ +-----+----+   +----------+ +----------+/
```
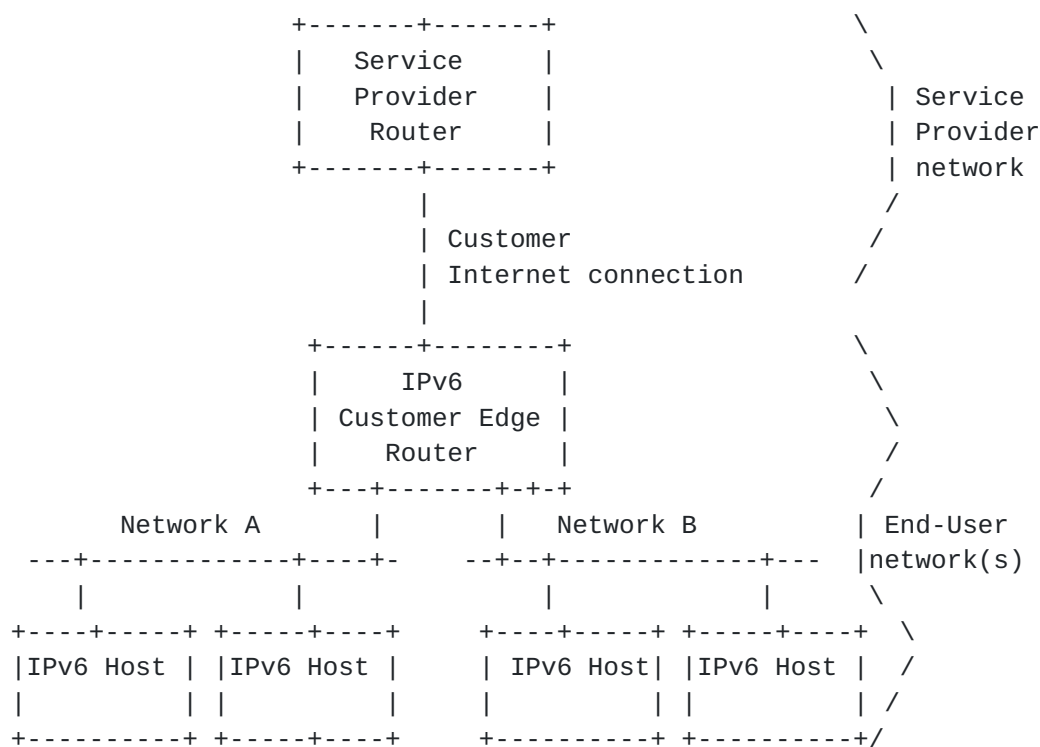
                              Figure 1

This architecture describes the:

o Basic capabilities of an IPv6 CE router

o Provisioning of the WAN interface connecting to the Service
   Provider

o Provisioning of the LAN interfaces

Unique Local IPv6 Unicast Addresses (ULA) [RFC4193] are used by hosts
communicating within the End-user Network; this is functionally
similar to RFC1918 addresses used within an IPv4 End-user Network.
The IPv6 CE router defaults to acting as the demarcation point
between two networks by providing a ULA boundary, a multicast zone
boundary and ingress and egress traffic filters.

For IPv6 multicast traffic the IPv6 CE router may act as an Multicast
Listener Discovery (MLD) proxy [RFC4605] and may support a dynamic
multicast routing protocol.

The IPv6 CE router may be manually configured in an arbitrary
topology with a dynamic routing protocol. Automatic provisioning and
configuration is described for a single IPv6 CE router only.

4. Requirements

4.1. General Requirements

The IPv6 CE router is responsible for implementing IPv6 routing; that
is, the IPv6 CE router must look up the IPv6 Destination address in
its routing table to decide to which interface it should send the
packet.

In this role, the IPv6 CE router is responsible for ensuring that
traffic using its ULA addressing does not go out the WAN interface,
and does not originate from the WAN interface.

G-1:   An IPv6 CE router is an IPv6 node according to the IPv6 Node
       Requirements [I-D.ietf-6man-node-req-bis] specification.

G-2:   The IPv6 CE router MUST implement ICMP according to [RFC4443].
       In particular point to point links MUST be handled as described
       in section 3.1 of [RFC4443].

G-3:   The IPv6 CE router MUST NOT forward any IPv6 traffic between
       its LAN Interface(s) and its WAN Interface until the router has
       successfully completed the IPv6 address acquisition process.

4.2. WAN Side Configuration

The IPv6 CE router will need to support connectivity to one or more
access network architectures. This document describes an IPv6 CE
router that is not specific to any particular architecture or Service
Provider, and supports all commonly used architectures.

IPv6 Neighbor Discovery and DHCPv6 protocols operate over any type of IPv6 supported link-layer and there is no need for a link-layer specific configuration protocol for IPv6 network layer configuration options as in e.g. PPP IPCP for IPv4. This section makes the assumption that the same mechanism will work for any link-layer, be it Ethernet, DOCSIS, PPP or others.

WAN side requirements:

W-1:   When the router is attached to the WAN interface link it MUST act as an IPv6 host for the purposes of stateless or stateful interface address assignment ([RFC4862] / [RFC3315]).

W-2:   The IPv6 CE router MUST generate a link-local address and finish Duplicate Address Detection according to [RFC4862] prior to sending any Router Solicitations on the interface. The source address used in the subsequent Router Solicitation MUST be the link-local address on the WAN interface.

W-3:   Absent of other routing information the IPv6 CE router MUST use Router Discovery as specified in [RFC4861] to discover a default router(s) and install default route(s) in its routing table with the discovered router's address as the next-hop.

W-4:   The router MUST act as a requesting router for the purposes of DHCPv6 prefix delegation ([RFC3633]).

W-5:   DHCPv6 address assignment (IA_NA) and DHCPv6 prefix delegation (IA_PD) SHOULD be done as a single DHCPv6 session.

W-6:   The IPv6 CE router MUST use a persistent DUID for DHCPv6 messages. The DUID MUST NOT change between network interface resets or IPv6 CE router reboot.

Link-layer requirements:

WLL-1: If the WAN interface supports Ethernet encapsulation, then the IPv6 CE router MUST support IPv6 over Ethernet [RFC2464].

WLL-2: If the WAN interface supports PPP encapsulation the IPv6 CE router MUST support IPv6 over PPP [RFC5072].

WLL-3: If the WAN interface supports PPP encapsulation, in a dual-stack environment with IPCP and IPV6CP running over one PPP logical channel, the NCPs MUST be treated as independent of each other and start and terminate independently.

Address assignment requirements:

WAA-1:   The IPv6 CE router MUST support SLAAC [RFC4862].

WAA-2:   The IPv6 CE router MUST follow the recommendation in
         [I-D.ietf-6man-ipv6-subnet-model] and in particular the
         handling of the L-flag in the Router Advertisement Prefix
         Information Option.

WAA-3:   The IPv6 CE router MUST support DHCPv6 [RFC3315] client
         behavior.

WAA-4:   The IPv6 CE router MUST be able to support the following
         DHCPv6 options: IA_NA, Reconfigure Accept [RFC3315],
         DNS_SERVERS [RFC3646].

WAA-5:   The IPv6 CE router SHOULD support the DHCPv6 SNTP option
         [RFC4075] and the Information Refresh Time Option [RFC4242].

WAA-6:   If the IPv6 CE router receives an RA message (described in
         [RFC4861]) with the M-flag set to 1, the IPv6 CE router MUST
         do DHCPv6 address assignment (request an IA_NA option).

WAA-7:   If the IPv6 CE router is unable to assign address(es) through
         SLAAC it MAY do DHCPv6 address assignment (request an IA_NA)
         even if the M-flag is set to 0.

WAA-8:   If the IPv6 CE router does not acquire global IPv6
         address(es) from either SLAAC or DHCPv6, then it MUST create
         global IPv6 address(es) from its delegated prefix(es) and
         configure those on one of its internal virtual network
         interfaces.

WAA-9:   As a router the IPv6 CE router MUST follow the weak host
         model [RFC1122]. When originating packets out an interface
         it will use a source address from another of its interfaces
         if the outgoing interface does not have an address of
         suitable scope.

   Prefix Delegation requirements:

WPD-1:   The IPv6 CE router MUST support DHCPv6 prefix delegation
         requesting router behavior as specified in [RFC3633] (IA_PD
         option).

WPD-2:   The IPv6 CE router MAY indicate as a hint to the delegating
         router the size of the prefix it requires. If so, it MUST
         ask for a prefix large enough to assign one /64 for each of
         its interfaces rounded up to the nearest nibble and MUST be
         configurable to ask for more.

WPD-3: The IPv6 CE router MUST be prepared to accept a delegated
        prefix size different from what is given in the hint. If the
        delegated prefix is too small to address all of its
        interfaces, the IPv6 CE router SHOULD log a system management
        error.

WPD-4: The IPv6 CE router MUST always initiate DHCPv6 prefix
        delegation, regardless of the M and O-flags in a received
        Router Advertisement message.

WPD-5: If the IPv6 CE Router initiates DHCPv6 before receiving a
        Router Advertisement it MUST also request an IA_NA option in
        DHCPv6.

WPD-6: If the delegated prefix(es) are aggregate route(s) of
        multiple, more-specific routes, the IPv6 CE router MUST
        discard packets that match the aggregate route(s), but not
        any of the more-specific routes. In other words, the next-
        hop for the aggregate route(s) should be the null
        destination. This is necessary to prevent forwarding loops
        when some addresses covered by the aggregate are not
        reachable [RFC4632].

        (a) The IPv6 CE router SHOULD send an ICMPv6 Destination
             Unreachable according to section 3.1 [RFC4443] back to
             the source of the packet, if the packet is to be dropped
             due to this rule.

WPD-7: If the IPv6 CE router requests both an IA_NA and an IA_PD in
        DHCPv6, it MUST accept an IA_PD in DHCPv6 Advertise/Reply
        messages, even if the message does not contain any addresses.

WPD-8: By default an IPv6 CE router MUST NOT initiate any dynamic
        routing protocol on its WAN interface.

4.3. LAN Side Configuration

The IPv6 CE router distributes configuration information obtained
during WAN interface provisioning to IPv6 hosts and assists IPv6
hosts in obtaining IPv6 addresses. It also supports connectivity of
these devices in the absence of any working WAN interface.

An IPv6 CE router is expected to support an IPv6 end-user network and
IPv6 hosts that exhibit the following characteristics:

1. Link-local addresses are insufficient for allowing IPv6
    applications to communicate with each other in the end-user
    network. The IPv6 CE router will need to enable this

communication by providing globally-scoped unicast addresses or
ULAs [RFC4193] whether or not WAN connectivity exists.

2. IPv6 hosts should be capable of using SLAAC and may be capable of
    using DHCPv6 for acquiring their addresses.

3. IPv6 hosts may use DHCPv6 for other configuration information,
    such as the DNS_SERVERS option for acquiring DNS information.

Unless otherwise specified, the following requirements apply to the
IPv6 CE router's LAN interfaces only.

Requirements:

L-1:   The IPv6 CE router MUST support ULA addressing [RFC4193].

L-2:   The IPv6 CE router MUST have a ULA prefix that it maintains
        consistently across reboots.

L-3:   The value of the ULA prefix SHOULD be user configurable.

L-4:   By default the IPv6 CE router MUST act as a site border router
        according to section 4.3 of [RFC4193] and filter packets with
        Local IPv6 source or destination addresses accordingly.

L-5:   The IPv6 CE router MUST support router behavior according to
        Neighbor Discovery for IPv6 [RFC4861].

L-6:   The IPv6 CE router MUST assign a separate /64 from its
        delegated prefix(es) (and ULA prefix if configured to provide
        ULA addressing) for each of its LAN interfaces.

L-7:   The IPv6 CE router MUST make each LAN interface an advertising
        interface according to [RFC4861].

L-8:   In Router Advertisements messages, the Prefix Information
        Option's A and L-flags MUST be set to 1 by default.

L-9:   The A and L-flags setting SHOULD be user configurable.

L-10:   The IPv6 CE router MUST support a DHCPv6 server capable of
        IPv6 address assignment according to [RFC3315] OR a stateless
        DHCPv6 server according to [RFC3736] on its LAN interfaces.

L-11:   Unless the IPv6 CE router is configured to support the DHCPv6
        IA_NA option, it SHOULD set M=0 and O=1 in its Router
        Advertisement messages [RFC4861].

L-12:   The IPv6 CE router MUST support providing DNS information in
        the DHCPv6 DNS_SERVERS option [RFC3646].

L-13:   The IPv6 CE router SHOULD make available a subset of DHCPv6
        options (as listed in section 5.3 of [RFC3736]) received from
        the DHCPv6 client on its WAN interface to its LAN side DHCPv6
        server.

L-14:   If the delegated prefix changes, i.e. the current prefix is
        replaced with a new prefix without any overlapping time
        period, then the IPv6 CE router MUST immediately advertise the
        old prefix with a preferred lifetime of 0 and a valid lifetime
        of 2 hours (which must be decremented in real time) in a
        Router Advertisement message.

L-15:   The IPv6 CE router MUST send an ICMP Destination Unreachable
        Message, code 5 (Source address failed ingress/egress policy)
        for packets forwarded to it using an address from a prefix
        which has been deprecated.

L-16:   If the IPv6 CE router loses its default route it SHOULD stop
        advertising itself as a default router. I.e. by setting the
        "Router Lifetime" field to 0 in subsequent Router
        Advertisement messages.

4.4. Security Considerations

It is considered a best practice to filter obviously malicious
traffic (e.g. spoofed packets, "martian" addresses, etc.). Thus, the
IPv6 CE router ought to support basic stateless egress and ingress
filters. The CE router is also expected to offer mechanisms to
filter traffic entering the customer network; however, the method by
which vendors implement configurable packet filtering is beyond the
scope of this document.

Security requirements:

S-1:   The IPv6 CE router SHOULD support
        [I-D.ietf-v6ops-cpe-simple-security]. In particular, the IPv6
        CE router SHOULD support functionality sufficient for
        implementing the set of recommendations in
        [I-D.ietf-v6ops-cpe-simple-security] section 4. Ths document
        takes no position on whether such functionality is enabled by
        default or mechanisms by which users would configure it.

S-2:   The IPv6 CE router MUST support ingress filtering in accordance
       with [RFC2827] (BCP 38)

5. Acknowledgements

Thanks to the following people (in alphabetical order) for their
guidance and feedback:

Mikael Abrahamsson, Merete Asak, Scott Beuker, Mohamed Boucadair, Rex
Bullinger, Brian Carpenter, Remi Denis-Courmont, Gert Doering, Alain
Durand, Katsunori Fukuoka, Tony Hain, Thomas Herbst, Kevin Johns,
Stephen Kramer, Victor Kuarsingh, Francois-Xavier Le Bail, David
Miles, Shin Miyakawa, Jean-Francois Mule, Michael Newbery, Carlos
Pignataro, John Pomeroy, Antonio Querubin, Teemu Savolainen, Matt
Schmitt, Hiroki Sato, Mark Townsley, Bernie Volz, James Woodyatt, Dan
Wing and Cor Zwart

This draft is based in part on CableLabs' eRouter specification. The
authors wish to acknowledge the additional contributors from the
eRouter team:

Ben Bekele, Amol Bhagwat, Ralph Brown, Eduardo Cardona, Margo Dolas,
Toerless Eckert, Doc Evans, Roger Fish, Michelle Kuska, Diego
Mazzola, John McQueen, Harsh Parandekar, Michael Patrick, Saifur
Rahman, Lakshmi Raman, Ryan Ross, Ron da Silva, Madhu Sudan, Dan
Torbet and Greg White

6. Contributors

The following people have participated as co-authors or provided
substantial contributions to this document: Ralph Droms, Kirk
Erichsen, Fred Baker, Jason Weil, Lee Howard, Jean-Francois Tremblay,
Yiu Lee, John Jason Brzozowski and Heather Kirksey.

7. IANA Considerations

This memo includes no request to IANA.

8. References

8.1. Normative References

[I-D.ietf-6man-ipv6-subnet-model]
          Singh, H., Beebee, W., and E. Nordmark, "IPv6 Subnet

                 Model: the Relationship between Links and Subnet
                 Prefixes", draft-ietf-6man-ipv6-subnet-model-12 (work in
                 progress), April 2010.

      [I-D.ietf-6man-node-req-bis]
                 Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node
                 Requirements RFC 4294-bis",
                 draft-ietf-6man-node-req-bis-05 (work in progress),
                 July 2010.

      [I-D.ietf-v6ops-cpe-simple-security]
                 Woodyatt, J., "Recommended Simple Security Capabilities in
                 Customer Premises Equipment for Providing Residential IPv6
                 Internet Service", draft-ietf-v6ops-cpe-simple-security-12
                 (work in progress), June 2010.

      [RFC1122] Braden, R., "Requirements for Internet Hosts -
                 Communication Layers", STD 3, RFC 1122, October 1989.

      [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
                 Requirement Levels", BCP 14, RFC 2119, March 1997.

      [RFC2464] Crawford, M., "Transmission of IPv6 Packets over Ethernet
                 Networks", RFC 2464, December 1998.

      [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering:
                 Defeating Denial of Service Attacks which employ IP Source
                 Address Spoofing", BCP 38, RFC 2827, May 2000.

      [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C.,
                 and M. Carney, "Dynamic Host Configuration Protocol for
                 IPv6 (DHCPv6)", RFC 3315, July 2003.

      [RFC3633] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic
                 Host Configuration Protocol (DHCP) version 6", RFC 3633,
                 December 2003.

      [RFC3646] Droms, R., "DNS Configuration options for Dynamic Host
                 Configuration Protocol for IPv6 (DHCPv6)", RFC 3646,
                 December 2003.

      [RFC3736] Droms, R., "Stateless Dynamic Host Configuration Protocol
                 (DHCP) Service for IPv6", RFC 3736, April 2004.

      [RFC4075] Kalusivalingam, V., "Simple Network Time Protocol (SNTP)
                 Configuration Option for DHCPv6", RFC 4075, May 2005.

      [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast

                     Addresses", RFC 4193, October 2005.

   [RFC4242]   Venaas, S., Chown, T., and B. Volz, "Information Refresh
               Time Option for Dynamic Host Configuration Protocol for
               IPv6 (DHCPv6)", RFC 4242, November 2005.

   [RFC4443]   Conta, A., Deering, S., and M. Gupta, "Internet Control
               Message Protocol (ICMPv6) for the Internet Protocol
               Version 6 (IPv6) Specification", RFC 4443, March 2006.

   [RFC4605]   Fenner, B., He, H., Haberman, B., and H. Sandick,
               "Internet Group Management Protocol (IGMP) / Multicast
               Listener Discovery (MLD)-Based Multicast Forwarding
               ("IGMP/MLD Proxying")", RFC 4605, August 2006.

   [RFC4632]   Fuller, V. and T. Li, "Classless Inter-domain Routing
               (CIDR): The Internet Address Assignment and Aggregation
               Plan", BCP 122, RFC 4632, August 2006.

   [RFC4779]   Asadullah, S., Ahmed, A., Popoviciu, C., Savola, P., and
               J. Palet, "ISP IPv6 Deployment Scenarios in Broadband
               Access Networks", RFC 4779, January 2007.

   [RFC4861]   Narten, T., Nordmark, E., Simpson, W., and H. Soliman,
               "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861,
               September 2007.

   [RFC4862]   Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless
               Address Autoconfiguration", RFC 4862, September 2007.

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

   [RFC5072]   S.Varada, Haskins, D., and E. Allen, "IP Version 6 over
               PPP", RFC 5072, September 2007.

8.2. Informative References

   [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-09 (work in progress),
               May 2010.

   [UPnP-IGD]
               UPnP Forum, "Universal Plug and Play (UPnP) Internet
               Gateway Device (IGD)", November 2001,

<http://www.upnp.org/standardizeddcps/igd.asp>.

Authors' Addresses

   Hemant Singh
   Cisco Systems, Inc.
   1414 Massachusetts Ave.
   Boxborough, MA 01719
   USA

   Phone: +1 978 936 1622
   Email: shemant@cisco.com
   URI: http://www.cisco.com/

   Wes Beebee
   Cisco Systems, Inc.
   1414 Massachusetts Ave.
   Boxborough, MA 01719
   USA

   Phone: +1 978 936 2030
   Email: wbeebee@cisco.com
   URI: http://www.cisco.com/

   Chris Donley
   CableLabs
   858 Coal Creek Circle
   Louisville, CO 80027
   USA

   Email: c.donley@cablelabs.com

   Barbara Stark
   AT&T
   725 W Peachtree St
   Atlanta, GA 30308
   USA

   Email: barbara.stark@att.com

Ole Troan (editor)
Cisco Systems, Inc.
Veversmauet 8
N-5017 BERGEN,
Norway

Email: ot@cisco.com