            IPv6 Multihoming without Network Address Translation
                  draft-troan-multihoming-without-nat66-01

Abstract

   Network Address and Port Translation (NAPT) works well for conserving
   global addresses and addressing multihoming requirements, because an
   IPv4 NAPT router implements three functions: source address
   selection, next-hop resolution and optionally DNS resolution. For
   IPv6 hosts one approach could be the use of IPv6 NAT. However, NAT
   should be avoided, if at all possible, to permit transparent host-to-
   host connectivity. In this document, we analyze the use cases of
   multihoming. We also describe functional requirements for
   multihoming without the use of NAT in IPv6 for hosts and small IPv6
   networks that would otherwise be unable to meet minimum IPv6
   allocation criteria .

Table of Contents

1. Introduction

    IPv6 provides enough globally unique addresses to permit every
    conceivable host on the Internet to be uniquely addressed without the
    requirement for Network Address Port Translation (NAPT [RFC3022])
    offering a renaissance in host-to-host transparent connectivity.

    Unfortunately, this may not be possible due to the necessity of NAT
    even in IPv6, because of multihoming.

    Multihoming is a blanket term to describe a host or small network
    that is connected to more than one upstream network. Whenever a host
    or small network (which does not meet minimum IPv6 allocation
    criteria) is connected to multiple upstream networks IPv6 addressing
    is assigned by each respective service provider resulting in hosts
    with more than one active IPv6 address. As each service provided is
    allocated a different address space from its Internet Registry, it
    in-turn assigns a different address space to the end-user network or
    host. For example, a remote access user may use a VPN to
    simultaneously connect to a remote network and retain a default route
    to the Internet for other purposes.

    In IPv4 a common solution to the multihoming problem is to employ
    NAPT on a border router and use private address space for individual
    host addressing. The use of NAPT allows hosts to have exactly one IP
    address visible on the public network and the combination of NAPT
    with provider-specific outside addresses (one for each uplink) and
    destination-based routing insulates a host from the impacts of
    multiple upstream networks. The border router may also implement a
    DNS cache or DNS policy to resolve address queries from hosts.

    It is our goal to avoid the IPv6 equivalent of NAT. To reach this
    goal, mechanisms are needed for end-user hosts to have multiple
    address assignments and resolve issues such as which address to use
    for sourcing traffic to which destination:

    o If multiple routers exist on a single link the host must
       appropriately select next-hop for each connected network. Routing
       protocols that would normally be employed for router-to-router
       network advertisement seem inappropriate for use by individual
       hosts.

    o Source address selection also becomes difficult whenever a host
       has more than one address within the same address scope. Current
       address selection criteria may result in hosts using an arbitrary
       or random address when sourcing upstream traffic. Unfortunately,
       for the host, the appropriate source address is a function of the
       upstream network for which the packet is bound for. If an

upstream service provider uses IP anti-spoofing or uRPF, it is
conceivable that the packets that have inappropriate source
address for the upstream network would never reach their
destination.

o In a multihomed environment, different DNS scopes or partitions
   may exist in each independent upstream network. A DNS query sent
   to an arbitrary upstream resolver may result in incorrect or
   poisoned responses.

In short, while IPv6 facilitates hosts having more than one address
in the same address scope, the application of this causes significant
issues for a host from routing, source address selection and DNS
resolution perspectives. A possible consequence of assigning a host
multiple identical-scoped addresses is severely impaired IP
connectivity.

If a host connects to a network behind an IPv4 NAPT, the host has one
private address in the local network. There is no confusion. The
NAT becomes the gateway of the host and forwards the packet to an
appropriate network when it is multihomed. It also operates a DNS
cache server, which receives all DNS inquires, and gives a correct
answer to the host.

In this document, we identify the functions present in multihomed
IPv4 NAPT environments and propose requirements that address
multihomed IPv6 environments without using IPv6 NAT.

2. Terminology

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].

NAT66 or IPv6 NAT   The terms "NAT66" and "IPv6 NAT" refer to
                      [I-D.mrw-behave-nat66].

NAPT   Network Address Port Translation as described
                      in [RFC3022]. In other contexts, NAPT is often
                      pronounced "NAT" or written as "NAT".

Multihomed with multi-prefix (MHMP) A host implementation which
                      supports the mechanisms described in this
                      document. Namely source address selection
                      policy, next-hop selection and DNS selection
                      policy.

3. IPv6 multihomed network scenarios

   In this section, we classify three scenarios of the multihoming
   environment.

3.1. Classification of network scenarios for multihomed host

   Scenario 1:

   In this scenario, two or more routers are present on a single link
   shared with the host(s). Each router is in turn connected to a
   different service provider network, which provides independent
   address assignment and DNS resolvers. A host in this environment
   would be offered multiple prefixes and DNS resolvers advertised from
   the two different routers.

```
                                +------+       _____
                                |      |      /           \
                           +---| rtr1 |=====/   network    \
                           |   |      |     \      1       /
          +------+         |   +------+      _____/
          |      |         |
          | host |-----+
          |      |         |
          +------+         |   +------+       _____
                           |   |      |      /           \
                           +---| rtr2 |=====/   network    \
                               |      |     \      2       /
                               +------+      _____/
```

             Figure 1: single uplink, multiple next-hop, multiple prefix
                                  (Scenario 1)

   Figure 1 illustrates the host connecting to rtr1 and rtr2 via a
   shared link. Networks 1 and 2 are reachable via rtr1 and rtr2
   respectively. When the host sends packets to network 1, the next-hop
   to network 1 is rtr1. Similarly, rtr2 is the next-hop to network 2.

   - e.g., broadband service (Internet, VoIP, IPTV, etc.)

   Scenario 2:

   In this scenario, a single gateway router connects the host to two or
   more upstream service provider networks. This gateway router would
   receive prefix delegations from each independent service provider
   network and a different set of DNS resolvers. The gateway in turn
   advertises the provider prefixes to the host, and for DNS, may either

act as a lightweight DNS resolver/cache or may advertise the complete
set of service provider DNS resolvers to the hosts.

```
                                      +------+      _____
                                      |      |     /            \
                                 +----| rtr1 |====/   network    \
                                 |    |      |    \      1       /
           +------+    +-----+   |    +------+     _____ _/
           |      |    |     |   |
           | host |----|  GW |---+
           |      |    | rtr |   |
           +------+    +-----+   |    +------+      _____
                                 |    |      |     /            \
                                 +----| rtr2 |====/   network    \
                                      |      |    \      2       /
                                      +------+     _____ _/
```
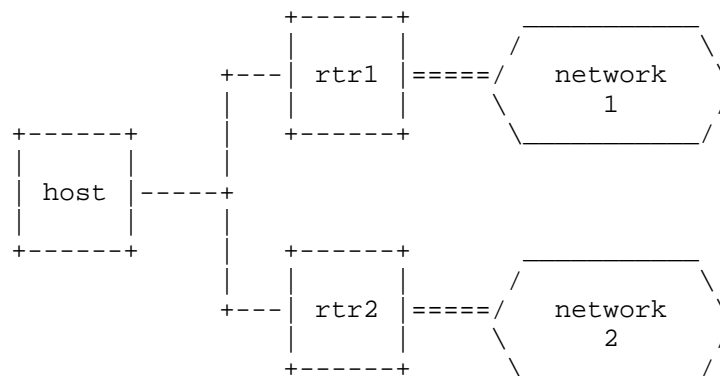
               Figure 2: single uplink, single next-hop, multiple prefix
                                (Scenario 2)

Figure 2 illustrates the host connected to GW rtr. GW rtr connects
to networks 1 and 2 via rtr1 and rtr2, respectively. When the host
sends packets to either network 1 or 2, the next-hop is GW rtr. When
the packets are sent to network 1 (network 2), GW rtr forwards the
packets to rtr1 (rtr2).

- e.g, Internet + VPN/ASP

Scenario 3:

In this scenario, a host has more than one active interfaces that
connects to different routers and service provider networks. Each
router provides the host with a different address prefix and set of
DNS resolvers, resulting in a host with a unique address per link/
interface.

```
                +------+     +------+      _____
                |      |     |      |     /            \
                |      |-----| rtr1 |====/   network    \
                |      |     |      |    \      1        /
                |      |     +------+     _____/
                |      |
                | host |
                |      |
                |      |     +------+      _____
                |      |     |      |     /            \
                |      |=====| rtr2 |====/   network    \
                |      |     |      |    \      2        /
                +------+     +------+     _____/
```
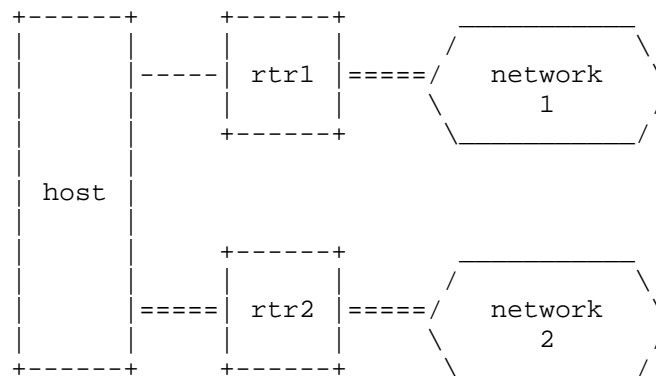
                Figure 3: Multiple uplink, multiple next-hop, multiple prefix
                                (Scenario 3)

   Figure 3 illustrates the host connecting to rtr1 and rtr2 via a
   direct connection or a virtual link. When the host sends packets
   network 1, the next-hop to network 1 is rtr1. Similarly, rtr2 is the
   next-hop to network 2.

   - e.g., Mobile Wifi + 3G, ISP A + ISP B

3.2. Multihomed network environment

   In an IPv6 multihomed network, a host is assigned two or more IPv6
   addresses and DNS resolvers from independent service provider
   networks. When this multihomed host attempts to connect with other
   hosts, it may incorrectly resolve the next-hop router, use an
   inappropriate source address, or use a DNS response from an incorrect
   service provider that may result in impaired IP connectivity.

   Multihomed networks in IPv4 have been commonly implemented through
   the use of a gateway router with NAPT function (scenario 2 with
   NAPT). An analysis of the current IPv4 NAPT and DNS functions within
   the gateway router should provide a baseline set of requirements for
   IPv6 multihomed environments. A destination prefix/route is often
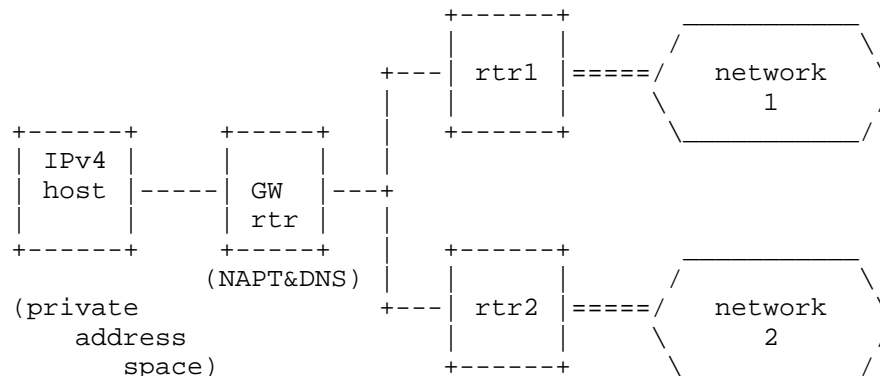   used on the gateway router to separate traffic between the networks.

```
                              +------+       _____
                              |      |      /           \
                        +---| rtr1 |=====/   network   \
                        |   |      |     \      1      /
    +------+   +-----+  |   +------+      _____/
    | IPv4 |   |     |  |   |
    | host |-----| GW   |---+
    |      |   | rtr  |  |
    +------+   +-----+  |   +------+       _____
             (NAPT&DNS) |   |      |      /           \
      (private          +---| rtr2 |=====/   network   \
         address            |      |     \      2      /
            space)          +------+      _____/
```

        Figure 4: IPv4 Multihomed environment with Gateway Router performing
                                      NAPT

3.3. Multihomed Problem Statement

   A multihomed IPv6 host has one or more assigned IPv6 addresses and
   DNS resolvers from each upstream service provider, resulting in the
   host having multiple valid IPv6 addresses and DNS resolvers. The
   host must be able to resolve the appropriate next-hop, the correct
   source address and DNS resolver to use based on the destination
   prefix. To prevent IP spoofing, operators will often implement IP
   filters and uRPF to discard traffic with an inappropriate source
   address, making it essential for the host to correctly resolve these
   three criteria before sourcing the first packet.

   IPv6 has mechanisms for the provision of multiple routers on a single
   link and multiple address assignments to a single host. However,
   when these mechanisms are applied to the three scenarios in
   Section 3.1 a number of connectivity issues are identified:

   Scenario 1:

   The host has been assigned an address from each router and recognizes
   both rtr1 and rtr2 as valid default routers (in the default routers
   list).

   o The source address selection policy on the host does not
      deterministically resolve a source address. Upstream uRPF or
      filter policies will discard traffic with source addresses that
      the operator did not assign.

   o The host will select one of the two routers as the active default
      router. No traffic is sent to the other router.

Scenario 2:

The host has been assigned two different addresses from the single
gateway router. The gateway router is the only default router on the
link.

o The source address selection policy on the host does not
   deterministically resolve a source address. Upstream uRPF or
   filter policies will discard traffic with source addresses that
   the operator did not assign.

o The gateway router does not have a mechanism for determining which
   traffic should be sent to which network. If the gateway router is
   implementing host functions (ie, processing RA) then two valid
   default routers may be recognized.

Scenario 3:

A host has two separate interfaces and on each interface a different
address is assigned. Each link has its own router.

o The host does not have enough information for determining which
   traffic should be sent to which upstream routers. The host will
   select one of the two routers as the active default router, and no
   traffic is sent to the other router.

o The default address selection rules select the address assigned to
   the outgoing interface as the source address. So, if a host has
   an appropriate routing table, an appropriate source address will
   be selected.

All scenarios:

o The host may use an incorrect DNS resolver for DNS queries.

4. Problem statement and analysis

The problems described in Section 3 can be classified into these
three types:

o Wrong source address selection

o Wrong next-hop selection

o Wrong DNS server selection

This section reviews the problem statements presented above and the

proposed functional requirements to resolve the issues without
employing IPv6 NAT.

4.1. Source address selection

   A multihomed IPv6 host will typically have different addresses
   assigned from each service provider either on the same link
   (scenarios 1 & 2) or different links (scenario 3). When the host
   wishes to send a packet to any given destination, the current source
   address selection rules [RFC3484] may not deterministically resolve
   the correct source address when the host addressing was via RA or
   DHCPv6. [I-D.ietf-6man-addr-select-sol] describes the use of the
   policy table [RFC3484] to resolve this problem, but there is no
   mechanism defined to disseminate the policy table information to a
   host. A proposal is in [I-D.fujisaki-dhc-addr-select-opt] to provide
   a DHCPv6 mechanism for host policy table management.

   Again, by employing DHCPv6, the server could restrict address
   assignment (of additional prefixes) only to hosts that support policy
   table management.

   Scenario 1: "Host" needs to support the solution for this problem

   Scenario 2: "Host" needs to support the solution for this problem

   Scenario 3: If "Host" support the next-hop selection solution, there
   is no need to support the address selection functionality on the
   host.

4.2. Next-hop selection

   A multihomed IPv6 host or gateway may have multiple uplinks to
   different service providers. Here each router would use Router
   Advertisements [RFC4861] for distributing default route/next-hop
   information to the host or gateway router.

   In this case, the host or gateway router may select any valid default
   router from the default routers list, resulting in traffic being sent
   to the wrong router and discarded by the upstream service provider.
   Using the above scenarios as an example, whenever the host wishes to
   reach a destination in network 2 and there is no connectivity between
   networks 1 and 2 (as is the case for a walled-garden or closed
   service), the host or gateway router does not know whether to forward
   traffic to rtr1 or rtr2 to reach a destination in network 2. The
   host or gateway router may choose rtr1 as the default router, and
   traffic fails to reach the destination server. The host or gateway
   router requires route information for each upstream service provider,
   but the use of a routing protocol between a host and router causes

both configuration and scaling issues. For IPv4 hosts, the gateway
router is often pre-configured with static route information or uses
of Classless Static Route Options [RFC3442] for DHCPv4. Extensions
to Router Advertisements through Default Router Preference and More-
Specific Routes [RFC4191] provides for link-specific preferences but
does not address per-host configuration in a multi-access topology
because of its reliance on Router Advertisements. A DHCPv6 option,
such as that in [I-D.dec-dhcpv6-route-option], is preferred for host-
specific configuration. By employing a DHCPv6 solution, a DHCPv6
server could restrict address assignment (of additional prefixes)
only to hosts that support more advanced next-hop and address
selection requirements.

Scenario 1: "Host" needs to support the solution for this problem

Scenario 2: "GW rtr" needs to support the solution for this problem

Scenario 3: "Host" needs to support the solution for this problem

4.3. DNS server selection

A multihomed IPv6 host or gateway router may be provided multiple DNS
resolvers through DHCPv6 or the experimental [RFC5006]. When the
host or gateway router sends a DNS query, it would normally choose
one of the available DNS resolvers for the query.

In the IPv6 gateway router scenario, the Broadband Forum [TR124]
required that the query be sent to all DNS resolvers, and the gateway
waits for the first reply. In IPv6, given our use of specific
destination-based policy for both routing and source address
selection, it is desirable to extend a policy-based concept to DNS
resolver selection. Doing so can minimize DNS resolver load and
avoid issues where DNS resolvers in different networks have
connectivity issues, or the DNS resolvers are not publicly
accessible. In the worst case, a DNS query may be unanswered if sent
towards an incorrect resolver, resulting in a lack of connectivity.

An IPv6 multihomed host or gateway router should have the ability to
select appropriate DNS resolvers for each service based on the domain
space for the destination, and each service should provide rules
specific to that network. [I-D.savolainen-mif-dns-server-selection]
proposes a solution for DNS server selection policy enforcement
solution with a DHCPv6 option.

Scenario 1: "Host" needs to support the solution for this problem

Scenario 2: "GW rtr" needs to support the solution for this problem

Scenario 3: "Host" needs to support the solution for this problem

5. Requirements

This section describes requirements that any solution multi-address and multi-uplink architectures need to meet.

5.1. End-to-End transparency

End-to-end transparency is a basic concept of the Internet. [RFC4966] states, "One of the major design goals for IPv6 is to restore the end-to-end transparency of the Internet. Therefore, because IPv6 is expected to remove the need for NATs and similar impediments to transparency, developers creating applications to work with IPv6 may be tempted to assume that the complex mechanisms employed by an application to work in a 'NATted' IPv4 environment are not required." The IPv6 multihoming solution SHOULD guarantee end-to-end transparency by avoiding IPv6 NAT.

5.2. Policy enforcement

The solution SHOULD have a function to enforce a policy on sites/ nodes. In particular, in a managed environment such as enterprise networks, an administrator has to control all nodes in his or her network.

The enforcement mechanisms should have:

o a function to distribute policies to nodes dynamically to update their behavior. When the network environment changes and the nodes' behavior has to be changed, a network administrator can modify the behavior of the nodes.

o a function to control every node centrally. A site administrator or a service provider could determine or could have an effect on the behavior at their users' hosts.

o a function to control node-specific behavior. Even when multiple nodes are on the same subnet, the mechanism should be able to provide a method for the network administrator to make nodes behave differently. For example, each node may have a different set of assigned prefixes. In such a case, the appropriate behavior may be different.

5.3. Scalability

   The solution will have to be able to manage a large number of sites/
   nodes. In services for residential users, provider edge devices have
   to manage thousands of sites. In such environments, sending packets
   periodically to each site may affect edge system performance.

6. Implementation approach

   As mentioned in Section 4, in the multi-prefix environment, we have
   three problems in source address selection, next-hop selection, and
   DNS resolver selection. In this section, possible solution
   mechanisms for each problem are introduced and evaluated against the
   requirements in Section 5.

6.1. Source address selection

   Possible solutions and their evaluation are summarized in
   [I-D.ietf-6man-addr-select-sol]. When those solutions are examined
   against the requirements in Section 5, the proactive approaches, such
   as the policy table distribution mechanism and the routing system
   assistance mechanism, are more appropriate in that they can propagate
   the network administrator's policy directly. The policy distribution
   mechanism has an advantage with regard to the host's protocol stack
   impact and the staticness of the assumed target network environment.

6.2. Next-hop selection

   As for the source address selection problem, both a policy-based
   approach and a non policy-based approach are possible with regard to
   the next-hop selection problem. Because of the same requirements,
   the policy propagation-based solution mechanism, whatever the policy,
   should be more appropriate.

   Routing information is a typical example of policy related to next-
   hop selection. If we assume source address-based routing at hosts or
   intermediate routers, the pairs of source prefixes and next-hops can
   be another example of next-hop selection policy.

   The routing information-based approach has a clear advantage in
   implementation and is already commonly used.

   The existing proposed or standardized routing information
   distribution mechanisms are routing protocols, such as RIPng and
   OSPFv3, the router advertisement (RA) extension option defined in
   [RFC4191], the DHCPv6 route information option proposed in
   [I-D.dec-dhcpv6-route-option], and the [TR069] standardized at BBF.

The RA-based mechanism has difficulty in per-host routing information
distribution. The dynamic routing protocols such as RIPng are not
usually used between the residential users and ISP networks because
of their scalability implications. The DHCPv6 mechanism does not
have these difficulties and has the advantages of its relaying
functionality. It is commonly used and is thus easy to deploy.

[TR069], mentioned above, is a possible solution mechanism for
routing information distribution to customer-premises equipment
(CPE). It assumes, however, IP reachability to the Auto
Configuration Server (ACS) is established. Therefore, if the CPE
requires routing information to reach the ACS, [TR069] cannot be used
to distribute this information.

6.3. DNS resolver selection

As in the above two problems, a policy-based approach and non policy-
based approach are possible. In a non policy-based approach, a host
or a home gateway router is assumed to send DNS queries to several
DNS servers at once or to select one of the available servers.

In the non policy-based approach, by making a query to a resolver in
a different service provider to that which hosts the service, a user
could be directed to unexpected IP address or receive an invalid
response, and thus cannot connect to the service provider's private
and legitimate service. For example, some DNS servers reply with
different answers depending on the source address of the DNS query,
which is sometimes called split-horizon. When the host mistakenly
makes a query to a different provider's DNS to resolve a FQDN of
another provider's private service, and the DNS resolver adopts the
split-horizon configuration, the queried server returns an IP address
of the non-private side of the service. Another problem with this
approach is that it causes unnecessary DNS traffic to the DNS
resolvers that are visible to the users.

The alternative of a policy-based approach is documented in
[I-D.savolainen-mif-dns-server-selection], where several pairs of DNS
resolver addresses and DNS domain suffixes are defined as part of a
policy and conveyed to hosts in a new DHCP option. In an environment
where there is a home gateway router, that router can act as a DNS
proxy, interpret this option and distribute DNS queries to the
appropriate DNS servers according to the policy.

7. Considerations for host without multi-prefix support

This section presents an alternative approach to mitigate the problem
in a multihomed network. This approach will help IPv6 hosts that are

not capable of the enhancements for the source address selection
policy, next-hop selection policy, and DNS selection policy described
in Section 6.

7.1. IPv6 NAT

In a typical IPv4 multihomed network deployment, IPv4 NAPT is
practically used and it can eventually avoid assigning multiple
addresses to the hosts and solve the next-hop selection problem. In
a similar fashion, IPv6 NAT can be used as a last resort for IPv6
multihomed network deployments where one needs to assign a single
IPv6 address to a host.

```
                                                        _____
                                                       /          \
                                              +---/  Internet  \
                            gateway router    |   \            /
     +------+    +--------------------+       |    _____/
     |      |    |      |             |  WAN1  +--+
     | host |-----|LAN|  Router |--------|
     |      |    |      |             |NAT|WAN2+--+
     +------+    +--------------------+       |     _____
                                              |    /          \
                                              +---/    ASP     \
                                                   \          /
                                                    _____/
```
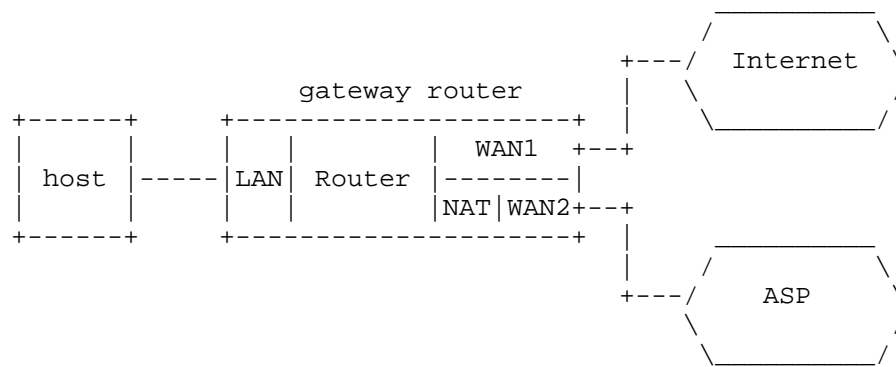
Figure 5: Legacy Host

The gateway router also has to support the two features, next-hop
selection and DNS server selection, shown in Section 6.

The implementation and issues of IPv6 NAT are out of the scope of
this document. They may be covered by another document under
discussion [I-D.mrw-behave-nat66].

7.2. Co-exisitence consideration

The above scenario relies on the assumption that only hosts without
multi-prefix support are connected to the GW rtr in scenario 2. To
allow the coexistence of non-MHMP hosts and MHMP hosts(i.e. hosts
supporting multi-prefix with the enhancements for the source address
selection), GW-rtr may need to treat those hosts separately.

An idea to achieve this is that GW-rtr identifies the hosts, and then
assigns single prefix to non-MHMP hosts and assigns multiple prefix
to MHMP hosts. In this case, GW-rtr can perform IPv6 NAT only for

the traffic from MHMP hosts if its source address is not appropriate.

Another idea is that GW-rtr assigns multiple prefix to the both hosts, and it performs IPv6 NAT for the traffic from non-MHMP hosts if its source address is not appropriate.

In scenario 1 and 3, the non-MHMP hosts can be placed behind the NAT box. In this case, non-MHMP host can access the service through the NAT box.

The implementation of identifying non-MHMP hosts and NAT policy is outside the scope of this document.

8. Security Considerations

This document does not define any new mechanisms. Each solution mechanisms should consider security risks independently. Security risks that occur as a result of combining solution mechanisms should be considered in another document.

9. IANA Considerations

This document has no IANA actions.

10. Contributors

The following people contributed to this document: Akiko Hattori, Arifumi Matsumoto, Frank Brockners, Fred Baker, Tomohiro Fujisaki, Jun-ya Kato, Shigeru Akiyama, Seiichi Morikawa, Mark Townsley, Wojciech Dec, Yasuo Kashimura, Yuji Yamazaki

11. References

11.1. Normative References

   [I-D.dec-dhcpv6-route-option]
             Dec, W. and R. Johnson, "DHCPv6 Route Option",
             draft-dec-dhcpv6-route-option-03 (work in progress),
             March 2010.

   [I-D.fujisaki-dhc-addr-select-opt]
             Fujisaki, T., Matsumoto, A., and R. Hiromi, "Distributing
             Address Selection Policy using DHCPv6",
             draft-fujisaki-dhc-addr-select-opt-09 (work in progress),

March 2010.

   [I-D.ietf-6man-addr-select-sol]
            Matsumoto, A., Fujisaki, T., and R. Hiromi, "Solution
            approaches for address-selection problems",
            draft-ietf-6man-addr-select-sol-03 (work in progress),
            March 2010.

   [I-D.mrw-behave-nat66]
            Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Address
            Translation (NAT66)", draft-mrw-behave-nat66-02 (work in
            progress), March 2009.

   [I-D.savolainen-mif-dns-server-selection]
            Savolainen, T., "DNS Server Selection on Multi-Homed
            Hosts", draft-savolainen-mif-dns-server-selection-02 (work
            in progress), February 2010.

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

   [RFC3484] Draves, R., "Default Address Selection for Internet
            Protocol version 6 (IPv6)", RFC 3484, February 2003.

   [RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and
            More-Specific Routes", RFC 4191, November 2005.

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

11.2. Informative References

   [RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network
            Address Translator (Traditional NAT)", RFC 3022,
            January 2001.

   [RFC3442] Lemon, T., Cheshire, S., and B. Volz, "The Classless
            Static Route Option for Dynamic Host Configuration
            Protocol (DHCP) version 4", RFC 3442, December 2002.

   [RFC4966] Aoun, C. and E. Davies, "Reasons to Move the Network
            Address Translator - Protocol Translator (NAT-PT) to
            Historic Status", RFC 4966, July 2007.

   [RFC5006] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli,
            "IPv6 Router Advertisement Option for DNS Configuration",
            RFC 5006, September 2007.

   [TR069]   The BroadBand Forum, "TR-069, CPE WAN Management Protocol
               v1.1, Version: Issue 1 Amendment 2", December 2007.

   [TR124]   The BroadBand Forum, "TR-124i2, Functional Requirements
               for Broadband Residential Gateway Devices (work in
               progress)", May 2010.

Authors' Addresses

   Ole Troan (editor)
   Cisco
   Bergen
   Norway

   Email: ot@cisco.com

   David Miles
   Alcatel-Lucent
   Melbourne
   Australia

   Email: david.miles@alcatel-lucent.com

   Satoru Matsushima
   SOFTBANK TELECOM Corp.
   Tokyo
   Japan

   Email: satoru.matsushima@tm.softbank.co.jp

   Tadahisa Okimoto
   NTT
   Tokyo
   Japan

   Email: t.okimoto@hco.ntt.co.jp

Dan Wing
Cisco
170 West Tasman Drive
San Jose
USA

Email: dwing@cisco.com