            Routing packets from hosts in a multi-prefix network
                   draft-ietf-6man-multi-homed-host-05

Abstract

   This document describes expected IPv6 host behavior in a network that
   has more than one prefix, each allocated by an upstream network that
   implements BCP 38 ingress filtering, when the host has multiple
   routers to choose from. It also applies to other scenarios such as
   the usage of stateful firewalls that effectively act as address-based
   filters. This host behavior may interact with source address
   selection in a given implementation, but logically follows it. Given
   that the network or host is, or appears to be, multihomed with
   multiple provider-allocated addresses, that the host has elected to
   use a source address in a given prefix, and that some but not all
   neighboring routers are advertising that prefix in their Router
   Advertisement Prefix Information Options, this document specifies to
   which router a host should present its transmission. It updates RFC
   4861.

Table of Contents

1. **Introduction and Applicability**

   This document describes the expected behavior of an IPv6 [RFC2460]
   host in a network that has more than one prefix, each allocated by an
   upstream network that implements BCP 38 [RFC2827] ingress filtering,
   and in which the host is presented with a choice of routers. It
   expects that the network will implement some form of egress routing,
   so that packets sent to a host outside the local network from a given
   ISP's prefix will go to that ISP. If the packet is sent to the wrong

egress, it is liable to be discarded by the BCP 38 filter. However, the mechanics of egress routing once the packet leaves the host are out of scope. The question here is how the host interacts with that network.

Various aspects of this issue, and possible solution approaches, are discussed in the document IPv6 Multihoming without Network Address Translation [RFC7157].

BCP 38 filtering by ISPs is not the only scenario where such behavior is valuable. Implementations that combine existing recommendations, such as [RFC6092] [RFC7084] can also result in such filtering. Another case is when the connections to the upstream networks include stateful firewalls, such that return packets in a stream will be discarded if they do not return via the firewall that created state for the outgoing packets. A similar cause of such discards is unicast reverse path forwarding (uRPF) [RFC3704].

In this document, the term "filter" is used for simplicity to cover all such cases. In any case, one cannot assume the host to be aware whether an ingress filter, a stateful firewall, or any other type of filter is in place. Therefore, the only consistent solution is to implement the features defined in this document.

Note that, apart from ensuring that a message with a given source address is given to a first-hop router that appears to know about the prefix in question, this specification is consistent with [RFC4861]. Nevertheless, implementers of Sections 5.2, 6.2.3, 6.3.4 and 8 of RFC 4861 will need to extend their implementations accordingly. This specification is fully consistent with [RFC6724] and implementers will need to add support for its Rule 5.5. Hosts that do not support these features may fail to communicate in the presence of filters as described above.

1.1. Host Model

It could be argued that the proposal of this document, which is to send messages using a source address in a given prefix to the router that advertised the prefix in its RA, is a form of [RFC1122]'s Strong ES (e.g. Host) Model, discussed in section 3.3.4.2 of that document. In short, [RFC1122] identifies two basic models, in which the "strong host" model models the host as a set of hosts in one chassis, each of which uses a single address on a single interface, and always both sends and receives on that interface, and the "weak host" model treats the host as one system with zero or more addresses on every interface, and capable of using any interface for any communication. As noted there, neither model is completely satisfactory. For example, a host with a link-local-only interface and a default route

pointing to the interface will necessarily send packets using any
address using that interface, and therefore be a de facto weak host.
If the router upstream from such a host implements BCP 38 Ingress
Filtering [RFC2827], such as by implementing uRPF on each interface,
the router might prevent communication by weak hosts.

```
              +-----------------+
              |                 |
              |     MIF Router  +---/--- other interfaces
              |                 |
              +---+---------+---+
                  |         | Two interfaces sharing a prefix
                --+-+--   --+-+--
                   |         |
                 +--+---------+--+
                 |   MIF Host    |
                 +---------------+
```
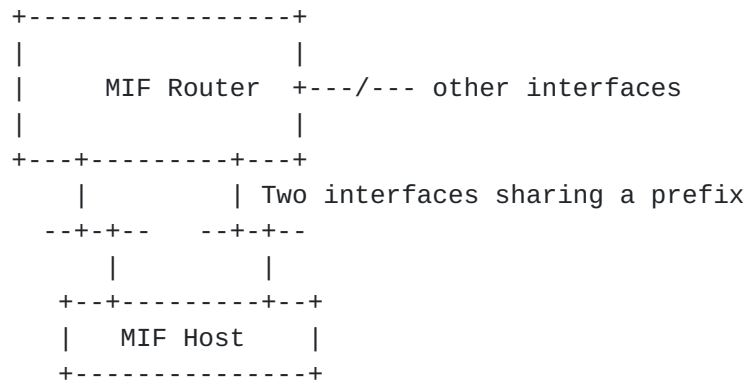
Figure 1: Hypothetical MIF interconnection

The proposal also differs slightly from [RFC1122]'s language of the
Strong Host Model. The statement is that the packet will go to the
router that advertised a given prefix, but doesn't state what
interface that might happen on. Hence, if the router is a MIF router
and is using the same prefix on two or more LANs shared by the host
(as in Figure 1), the host might use each of those LANs and meet the
requirement. The Strong Host Model is not stated in those terms, but
in terms of the interface used, and would find a MIF router quite
confusing:

   (A) A host [MUST] silently discard an incoming datagram whose
   destination address does not correspond to the physical interface
   through which it is received.

   (B) A host [MUST] restrict itself to sending (non-source- routed)
   IP datagrams only through the physical interface that corresponds
   to the IP source address of the datagrams.

However, comparing the presumptive route lookup mechanisms in each
model, this proposal is indeed most similar to the Strong Host Model,
as is any source/destination routing paradigm.

Strong: route(src IP addr, dest IP addr, TOS) -> gateway

Weak: route(dest IP addr, TOS) -> gateway, interface

In the hypothetical MIF model suggested in Figure 1, the address
fails to identify a single interface, but it does identify a single
gateway.

1.2. 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 [RFC2119].

2. Sending context expected by the host

2.1. Expectations the host has of the network

A host receives prefixes in a Router Advertisement [RFC4861], which
goes on to identify whether they are usable by SLAAC [RFC4862]
[RFC4941] [RFC7217]. When no prefixes are usable for SLAAC, the
Router Advertisement would normally signal the availability of DHCPv6
[RFC3315] and the host would use it to configure its addresses. In
the latter case (or if both SLAAC and DHCPv6 are used on the same
link for some reason) it will generally be the case that the
configured addresses match one of the prefixes advertised in a Router
Advertisement that are supposed to be on-link for that link.

The simplest multihomed network implementation in which a host makes
choices among routers might be a LAN with one or more hosts on it and
two or more routers, one for each upstream network, or a host that is
served by disjoint networks on separate interfaces. In such a
network, especially the latter, there is not necessarily a routing
protocol, and the two routers may not even know that the other is a
router as opposed to a host, or may be configured to ignore its
presence. One might expect that the routers may or may not receive
each other's RAs and form an address in the other router's prefix
(which is not per [RFC4862], but is implemented by some stub router
implementations). However, all hosts in such a network might be
expected to create an address in each prefix so advertised.

```
  +---------+   +---------+    +---------+    +---------+
  |  ISP    |   |  ISP    |    |  ISP    |    |  ISP    |
  +----+----+   +----+----+    +----+----+    +----+----+
       |             |              |              |
       |             |              |              |
  +----+----+   +----+----+    +----+----+    +----+----+
  | Router  |   | Router  |    | Router  |    | Router  |
  +----+----+   +----+----+    +----+----+    +----+----+
       |             |              |              |
       +------+------+              |   +--------+  |
              |                     +--+  Host  +--+
         +----+----+                   +--------+
         |  Host   |
         +---------+
         Common LAN Case            Disjoint LAN Case
        (Multihomed Network)        (Multihomed Host)
```
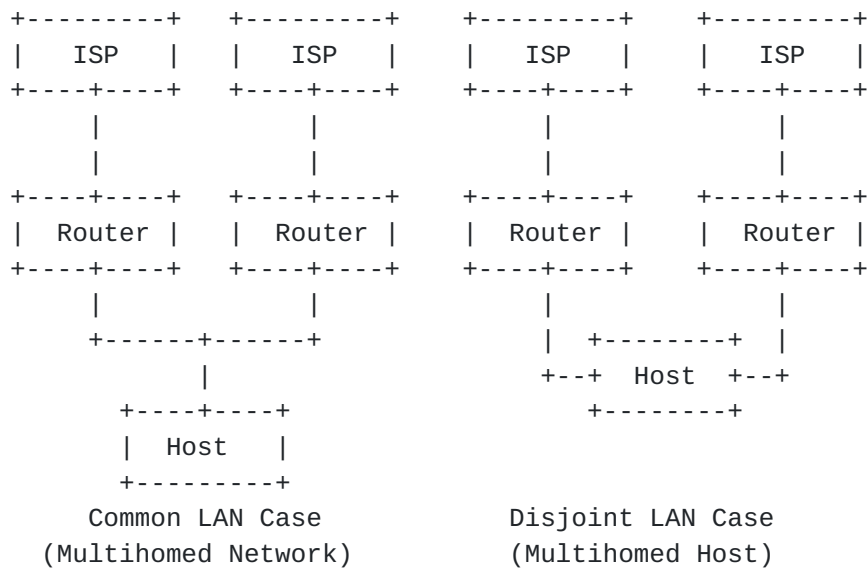
                   Figure 2: Two simple networks

   If there is no routing protocol among those routers, there is no
   mechanism by which packets can be deterministically forwarded between
   the routers (as described in BCP 84 [RFC3704]) in order to avoid
   filters. Even if there was routing, it would result in an indirect
   route, rather than a direct route originating with the host; this is
   not "wrong", but can be inefficient. Therefore the host would do
   well to select the appropriate router itself.

   Since the host derives fundamental default routing information from
   the Router Advertisement, this implies that, in any network with
   hosts using multiple prefixes, each prefix SHOULD be advertised via a
   Prefix Information Option (PIO) [RFC4861] by one of the attached
   routers, even if addresses are being assigned using DHCPv6. A router
   that advertises a prefix indicates that it is able to appropriately
   route packets with source addresses within that prefix, regardless of
   the setting of the L and A flags in the PIO.

   In some circumstances both L and A might be zero. If SLAAC is not
   wanted (A=0) and there is no reason to announce an on-link prefix
   (L=0), a PIO SHOULD be sent to inform hosts that the prefix is
   source-routed by the router in question. Although this does not
   violate the existing standard [RFC4861], such a PIO has not
   previously been common, and it is possible that existing host
   implementations simply ignore such a PIO or that a router
   implementation rejects such a PIO as a configuration error. Newer
   implementations that support this mechanism will need to be updated
   accordingly: a host SHOULD NOT ignore a PIO simply because both L and
   A flags are cleared; a router SHOULD be able to send such a PIO.

2.2. Expectations of multihomed networks

The direct implication of Section 2.1 is that, if the network uses a
routing protocol, the routing protocols used in multihomed networks
SHOULD implement source-prefix based egress routing, for example as
described in [I-D.ietf-rtgwg-dst-src-routing]. Network designs exist
that can usefully limit themselves to static routing (such as a
simple tree network), or may internally use no routers at all, such
as a single LAN with two CE routers, each of which leads to a
different upstream network.

3. Reasonable expectations of the host

3.1. Interpreting Router Advertisements

As described in [RFC4191] and [RFC4861], a Router Advertisement may
contain zero or more Prefix information Options (PIOs), or zero or
more Route Information Options (RIOs). In their original intent,
these indicate general information to a host: "the router whose
address is found in the source address field of this packet is one of
your default routers", "you might create an address in this prefix",
or "this router would be a good place to send traffic directed to a
given destination prefix". In a multi-homed network implementing
source/destination routing, the interpretation of default router or
an RIO has to be modified with the words "if the source address is in
one of the prefixes I advertise in a PIO". Additionally, the PIO
must be reinterpreted to also imply that the advertising router would
be a reasonable first hop for any packet using a source address in
any advertised prefix.

```
                                    +---------+  |
                         ( ISP A ) - +  Bob-A  +--+  +-----+
  +-------+                  /            +---------+  +--+     |
  |       |                 /                          |  |     |
  | Alice +--/--( The Internet )                       | Bob |
  |       |                 \                          |  |     |
  +-------+                  \            +---------+  +--+     |
                         ( ISP B ) - +  Bob-B  +--+  +-----+
                                    +---------+  |
```

                Figure 3: PIOs, RIOs, and Default Routes
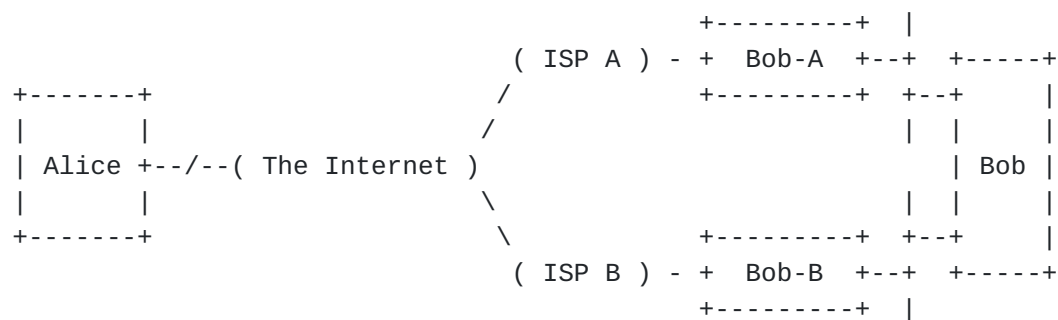
The implications bear consideration. Imagine, Figure 3, that hosts
Alice and Bob are in communication, Bob's network is multihomed, and
Bob's first hop routers are Bob-A (to upstream ISP A advertising
prefix PA) and Bob-B (to upstream network B and advertising prefix
PB). If Bob is responding to a message from Alice, his choice of
source address is forced to be the address Alice used as a

destination (which we may presume to have been in prefix PA). Hence,
Bob created or was assigned an address in PA, and can only reasonably
send traffic using it to Bob-A as a first hop router. If there were
several instances of Bob-A and one had advertised itself as a default
router or as having a route to Alice, that is the router Bob should
choose. If none of Bob-A have advertised that but Bob-B has, it is
irrelevant; Bob is using the address allocated in PA and courts a BCP
38 discard if he doesn't send the packet to Bob-A.

In the special case that Bob is initiating the conversation, an RIO
might, however, influence source address choice. Bob could
presumably use any address allocated to him, in this case his address
in PA or PB. If Bob-B has advertised an RIO for Alice's prefix and
Bob-A has not, Bob MAY take that fact into account in address
selection - choosing an address that would allow him to make use of
the RIO.

3.2. Default Router Selection

Default Router Selection is modified as follows: A host SHOULD select
default routers for each prefix it is assigned an address in.
Routers that have advertised the prefix in its Router Advertisement
message SHOULD be preferred over routers that do not advertise the
prefix.

As a result of doing so, when a host sends a packet using a source
address in one of those prefixes and has no history directing it
otherwise, it SHOULD send it to the indicated default router. In the
"simplest" network described in Section 2.1, that would get it to the
only router that is directly capable of getting it to the right ISP.
This will also apply in more complex networks, even when more than
one physical or virtual interface is involved.

In more complex cases, wherein routers advertise RAs for multiple
prefixes whether or not they have direct or isolated upstream
connectivity, the host is dependent on the routing system already.
If the host gives the packet to a router advertising its source
prefix, it should be able to depend on the router to do the right
thing.

3.3. Source Address Selection

There is an interaction with Default Address Selection [RFC6724]. A
host following the recommendation in the previous section will store
information about which next-hops advertised which prefixes. Rule
5.5 of RFC 6724 states that the source address used to send to a
given destination address should if possible be chosen from a prefix
known to be advertised by the next-hop router for that destination.

This selection rule would therefore be applicable in a host following
the recommendation in the previous section.

3.4. Redirects

There is potential for adverse interaction with any off-link Redirect
(Redirect for a destination that is not on-link) message sent by a
router in accordance with Section 8 of [RFC4861]. Hosts SHOULD apply
off-link redirects only for the specific pair of source and
destination addresses concerned, so the host's Destination Cache may
need to contain appropriate source-specific entries.

3.5. History

Some modern hosts maintain history, in terms of what has previously
worked or not worked for a given address or prefix and in some cases
the effective window and MSS values for TCP or other protocols. This
might include a next hop address for use when a packet is sent to the
indicated address.

When such a host makes a successful exchange with a remote
destination using a particular address pair, and the host has
previously received a PIO that matches the source address, then the
host SHOULD include the prefix in such history, whatever the setting
of the L and A flags in the PIO. On subsequent attempts to
communicate with that destination, if it has an address in that
prefix at that time, a host MAY use an address in the remembered
prefix for the session.

4. Residual issues

Consider a network where routers on a link run a routing protocol and
are configured with the same information. Thus, on each link all
routers advertise all prefixes on the link. The assumption that
packets will be forwarded to the appropriate egress by the local
routing system might cause at least one extra hop in the local
network (from the host to the wrong router, and from there to another
router on the same link).

In a slightly more complex situation such as the disjoint LAN case of
Figure 2, which happens to be one of the authors' home plus corporate
home-office configuration, the two upstream routers might be on
different LANs and therefore different subnets (e.g., the host is
itself multi-homed). In that case, there is no way for the "wrong"
router to detect the existence of the "right" router, or to route to
it.

In such a case it is particularly important that hosts take the
responsibility to memorize and select the best first-hop as described
in Section 3.

5. IANA Considerations

This memo asks the IANA for no new parameters.

6. Security Considerations

This document does not create any new security or privacy exposures.
It is intended to avoid connectivity issues in the presence of BCP 38
ingress filters or stateful firewalls combined with multihoming.

There might be a small privacy improvement, however: with the current
practice, a multihomed host that sends packets with the wrong address
to an upstream router or network discloses the prefix of one upstream
to the other upstream network. This practice reduces the probability
of that occurrence.

7. Acknowledgements

Comments were received from Jinmei Tatuya and Ole Troan, who have
suggested important text, plus Mikael Abrahamsson, Steven Barth,
Juliusz Chroboczek, Toerless Eckert, David Farmer, Dusan Mudric,
Tadahisa Okimoto, Pierre Pfister, Behcet Sarikaya, Mark Smith, Bob
Hinden, and James Woodyatt.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
           Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/
           RFC2119, March 1997,
           <http://www.rfc-editor.org/info/rfc2119>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6
           (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460,
           December 1998, <http://www.rfc-editor.org/info/rfc2460>.

[RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and
           More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191,
           November 2005, <http://www.rfc-editor.org/info/rfc4191>.

   [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman,
              "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861,
              DOI 10.17487/RFC4861, September 2007,
              <http://www.rfc-editor.org/info/rfc4861>.

   [RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown,
              "Default Address Selection for Internet Protocol Version 6
              (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012,
              <http://www.rfc-editor.org/info/rfc6724>.

8.2. Informative References

   [I-D.ietf-rtgwg-dst-src-routing]
              Lamparter, D., "Destination/Source Routing", draft-ietf-
              rtgwg-dst-src-routing-00 (work in progress), October 2015.

   [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts -
              Communication Layers", STD 3, RFC 1122, DOI 10.17487/
              RFC1122, October 1989,
              <http://www.rfc-editor.org/info/rfc1122>.

   [RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering:
              Defeating Denial of Service Attacks which employ IP Source
              Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827,
              May 2000, <http://www.rfc-editor.org/info/rfc2827>.

   [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins,
              C., and M. Carney, "Dynamic Host Configuration Protocol
              for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July
              2003, <http://www.rfc-editor.org/info/rfc3315>.

   [RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed
              Networks", BCP 84, RFC 3704, DOI 10.17487/RFC3704, March
              2004, <http://www.rfc-editor.org/info/rfc3704>.

   [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless
              Address Autoconfiguration", RFC 4862, DOI 10.17487/
              RFC4862, September 2007,
              <http://www.rfc-editor.org/info/rfc4862>.

   [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy
              Extensions for Stateless Address Autoconfiguration in
              IPv6", RFC 4941, DOI 10.17487/RFC4941, September 2007,
              <http://www.rfc-editor.org/info/rfc4941>.

   [RFC6092] Woodyatt, J., Ed., "Recommended Simple Security
              Capabilities in Customer Premises Equipment (CPE) for
              Providing Residential IPv6 Internet Service", RFC 6092,
              DOI 10.17487/RFC6092, January 2011,
              <http://www.rfc-editor.org/info/rfc6092>.

   [RFC7084] Singh, H., Beebee, W., Donley, C., and B. Stark, "Basic
              Requirements for IPv6 Customer Edge Routers", RFC 7084,
              DOI 10.17487/RFC7084, November 2013,
              <http://www.rfc-editor.org/info/rfc7084>.

   [RFC7157] Troan, O., Ed., Miles, D., Matsushima, S., Okimoto, T.,
              and D. Wing, "IPv6 Multihoming without Network Address
              Translation", RFC 7157, DOI 10.17487/RFC7157, March 2014,
              <http://www.rfc-editor.org/info/rfc7157>.

   [RFC7217] Gont, F., "A Method for Generating Semantically Opaque
              Interface Identifiers with IPv6 Stateless Address
              Autoconfiguration (SLAAC)", RFC 7217, DOI 10.17487/
              RFC7217, April 2014,
              <http://www.rfc-editor.org/info/rfc7217>.

Appendix A. Change Log (RFC Editor: please delete)

   Initial Version: 2015-08-05

   Version 01: Update text on PIOs, added text on Redirects, and
      clarified the concept of a "simple" network, 2015-08-13.

   Version 02: Clarifications after WG discussions, 2015-08-19.

   Version 03: More clarifications after more WG discussions,
      especially adding stateful firewalls, uRPF, and more precise
      discussion of RFC 4861, 2015-09-03.

   Version 04: Responds to various comments including

      * Questions regarding RFC 1122's strong and weak host models.
         This model is, strictly speaking, neither, but is most similar
         to the strong host model.

      * Some wording errors.

      * Requests for discussion of the handling of the RIO, PIO, and
         Default Router List in an RA.

   WG Versions 00-02: More clarifications after more WG discussions,
      2015-11-03.

      WG Version 03:   A final clarification re uRPF, 2015-12-15.

Authors' Addresses

    Fred Baker
    Cisco Systems
    Santa Barbara, California 93117
    USA

    Email: fred@cisco.com

    Brian Carpenter
    Department of Computer Science
    University of Auckland
    PB 92019
    Auckland 1142
    New Zealand

    Email: brian.e.carpenter@gmail.com