IPv6 Maintenance Working Group
Internet-Draft
Intended status: Informational
Expires: September 9, 2010

A. Matsumoto
T. Fujisaki
NTT
R. Hiromi
Intec Netcore

March 8, 2010

Considerations of address selection policy conflicts
draft-arifumi-6man-addr-select-conflict-02.txt

Abstract

This document examines how policy conflicts happen, and how to address the conflicts. After making it clear what kind of address selection policy should be necessary, we proposed how to merge the possibily conflicting policies for each of the destination address selection policy and source address selection policy.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

RFC 5220 [RFC5220] describes several cases where problems are caused by using multiple prefixes at hosts and sites. The address selection design team is working on this issue and summarizes their work in the considerations document [I-D.ietf-6man-addr-select-considerations]. Their solution mechanism is going to be an update mechanism of the address selection policy table at a host from the network. As one of the possible solutions, a DHCPv6 option [I-D.fujisaki-dhc-addr-select-opt] is proposed.

As mentioned in RFC 5220 [RFC5220], a host or a site belongs to multiple upstream networks in some environments. For example, a host with multiple interfaces, such as wireless and wired interfaces, can easily belong to multiple networks. A site may have connectivity to ISP and a corporate network through a VPN link.

In these cases, if two or more of the upstream networks want to control address selection behavior of his network's customer host, those address selection policies have to be merged at the host, and they may collide there.

Note that this document does not assume the address selection policy transmitted by an upstream network provider is in the form of RFC 3484 [RFC3484] policy table itself. Rather, it sorts out the motivation for address selection control in Section 2, and accordingly defines the form of policy. Then, this document describes how policy conflicts happen, and how they can be merged in the Section 4.

Some of the problems described in RFC 5220 are specific to and resulted from the address selection mechanism defined in RFC 3484 [RFC3484]. However, above mentioned policy collision is an intrinsic problem of address selection policy merging, and not specific to the RFC 3484 mechanism.

2. Motivations for Address Selection Control

As in RFC 5220, there are various motivations for network administrator to control address selection behavior of his customers' hosts. We can classify these policies into the following two groups.

- Source address selection behavior control:

"When accessing to PREFIX-1, use ADDRESS-1 as the source address." A lot of ISPs have this policy and they usually implement it by adopting ingress filtering to incoming packets from their customers. Another example is a multi-prefix network that makes use of multiple address blocks in the network and assigns multiple addresses/prefixes to its customers for different purpose, such as a prefix for the Internet access and a prefix for telephone call.

- Destination address selection behavior control:

  "When accessing to PREFIX-1 or PREFIX-2, prefer PREFIX-1 rather than PREFIX-2." This kind of control policy is usually intended for optimization of the customers' traffic. It is not usually intended for on-off switch manner of control, but rather control of preference degree. For example, this is useful when a destination site has both PREFIX-1 and PREFIX-2, and the network administrator knows connectivity to PREFIX-1 is better than PREFIX-2. The typical use case is IPv4 and IPv6 prioritization as mentioned in RFC 5220.

  On-off switch manner of destination address selection control is not in scope of RFC 3484 address selection behavior, but it can be implemented by some other means, such as routing table manipulation and DNS resolution.

  As it is intrinsiclly intended for optimization, it should not be used for any other purpose like security .

Here, PREFIX-* is used to denote both IPv4 and IPv6 prefixes. In the following part, policy conflict and its solution for these two kinds of control policy are discussed individually.

3. Conflicts and Solution Analysis

3.1. Source Address Selection

As mentioned above, source address selection policy have following meaning:

"When accessing to PREFIX-1, use ADDRESS-1 as the source address."

The upstream network that has this kind of policy usually assigns an address block that includes ADDRESS-1, and also provides reachability to the network that is specified by PREFIX-1.

Source address selection policy conflict can happen when different network have a policy for the same prefix. For example, in the

following figure, Network-1 have a policy: "To PREFIX-1, use ADDRESS-1", and Network-2: "To PREFIX-1 and PREFIX-2, use ADDRESS-2".

```
                 PREFIX-1 -----+    PREFIX-2
                       |        \     |
                       |         \    |
                 +-----+-----+  +-+---+-----+
                 | Network-1 |  | Network-2 |
                 +------+----+  +----+------+
                         \          /
                          \        /
                 ADDRESS-1 \      / ADDRESS-2
                        +---+----+---+
                        | Host/Site  |
                        +------------+
```

In this case, the solution is straightforward. As documented in RFC 3484, the destination address is determined before source address selection policy is used. Thus, the outgoing route, such as the next-hop node and the network interface, is determined by looking up the routing table at a host. In other words, the outgoing network that carries the packet to the destination is determined without the source address selection policy.

So, the bottom line is that the source address selection policy that matches routing table's behavior should be chosen. There is no point adopting the source address selection policy of a network where a packet does not go through.

In other words, if the routing table is fixed before the source address selection policy is fixed, then the source address selection policy should be implemented while avoiding contradiction with the routing table. If not, the routing table should be coordinated to match the source address selection policy.

In a case where a site is connected to the multiple ISPs, like the figure above, and receives policies from the ISPs and re-distribute policies to the downstream hosts, the hosts cannot know which ISP are chosen for transit to PREFIX-1. So, in this case, the entity who knows which way is chosen have to address the policy conflict.

For example, Network-1 and Network-2 advertise the following policy to the customers,

```
        Network-1: to PREFIX-1, use ADDRESS-1
        Network-2: to PREFIX-1, use ADDRESS-2
```

to PREFIX-2, use ADDRESS-2

The policy for PREFIX-1 is conflicted in this case. And when the routing table of the Host/Site is like below,

Destination	Gateway
PREFIX-1	Gateway Address of Network-1
PREFIX-2	Gateway Address of Network-2

the merging process should choose a policy from Network-1 for the conflicting PREFIX-1. By this process, the merged policy table will be:

to PREFIX-1, use ADDRESS-1
to PREFIX-2, use ADDRESS-2

3.2. Destination Address Selection

As mentioned in section 2, destination address selection policy have following meaning: "When accessing to a destination site that has PREFIX-1 and PREFIX-2, prefer PREFIX-1 rather than PREFIX-2." The upstream network that has this kind of policy should provides reachability to both networks that are specified by PREFIX-1 and PREFIX-2.

Destination address selection policy conflict can happen when a network has a policy that has inverse effect of another network's policy. That is, in the figure below, Network-1 prefers PREFIX-1 rather than PREFIX-2, and Network-2 prefers PREFIX-2 rather than PREFIX-1.

```
         bad   bad
PREFIX-1 ---- ---- PREFIX-2
      |      X      |
 good |     / \     | good
+-----+-----+ +-----+-----+
| Network-1 | | Network-2 |
+------+----+ +----+------+
        \         /
         \       /
ADDRESS-1 \     / ADDRESS-2
       +---+---+---+
       | Host/Site |
       +-----------+
```

In routing mechanism, a router advertises a route A to a certain destination, another advertises a route B to the same destination, and the receiver decides which route to take by looking at the cost of the routes and other information.

In destination address selection policy, a network advertises prefix A and the precedence degree. The destination address selection policy conflict happens when multiple entities provide policies for the same or the overlapping destination prefix with different costs. This is the same situation as the routing mechanism in that there can be multiple "routes" for the same destination.

Here, we can choose the better, that is, higher precedence "route" for the destination prefix, but there is no point if the route is not actually used by the routing mechainism. Even if we choose a policy for prefix A provided from Network-1, a packet destined for the prefix A does not always go through Network-1. This is what the routing mechianism of the host or the site router decides.

So, we propose to adopt the policy that is provided from the network the routing mechanism selected and thus a packet goes through, because the routing mechanism is already there and performs routing decisions by making use of the routing protocols metrics and also implementation dependent information.

For example, Network-1 router advertises the following policy to the customers:

```
to PREFIX-1, precedence 20
to PREFIX-2, precedence 10
```

Network-2 advertises:

```
to PREFIX-1, precedence 30
to PREFIX-2, precedence 40
```

And when the routing table is:

```
PREFIX-1 via Network-1
PREFIX-2 via Network-2
```

Then, the receiving host should have the following merged destination address selection policy:

      to PREFIX-1, precedence 20 via Network-1
      to PREFIX-2, precedence 40 via Network-2

4. Cenceptual Policy Processing Model

4.1. The Whole Picture of Policy Merging Process

The merging process in Section 3 describes how conflicting source address selection policies can be merged, and how conflicting destination address selection policies can be merged. The output of these merged process is not in the form of RFC 3484 policy table.

However, we can get these merged policies into the RFC 3484 policy table by the following processing of the merged policies. By processing the policies this way, we do not need to modify the form of RFC 3484 policy table.

```
           +-------------+ +-------------+
           |  Network-1  | |  Network-2  |
           +-------------+ +-------------+
                  |      \ /      |
                  |       X       |
                  v      v v      v
           +-------------+ +-------------+
           | Dst Address | | Src Address |
           | Policy Pool | | Policy Pool |
           +-------------+ +-------------+
                      |       |             +---------------+
                      |<------|<------------| Routing Table |
                      v       v             +---------------+
                 +-----------------+
                 | RFC3484 Default |
                 |  Policy Table   |
                 +-----------------+
                          |
                          v
           +-----------------------------+
           |     RFC 3484 Policy Table   |
           +-----------------------------+
```

In the figure above, Network-1 and Network-2 provide both destination address selection poilcies and source address selection policies. The policy receiver should keep the received policies as they are in policy pools. Also, the default policy table define in RFC 3484 has to be kept.

The conflicts in policy pools can be solved by looking up the routing table by the process of Section 3. The outputs from this process are injected into the RFC 3484 default policy table so that the injecting policies should least change the effects of the default policy.

This process should be performed every time the received policy changes. This is because the change cannot alway be made incrementally.

In the sub-sections below, the detailed merging process are described. The merging process begins with the merging of source address selection policies. The merging of destination address policies comes after that. And it ends with adjustment of the inserted entries' precedence and label values in accordance with the default policy table.

4.2. Example of Address Selection Policy Merging

```
                   ::ffff:0:0/96
                        ::/0 -----+
                          |        \
                          |         \
                    +-----+-----+ +--+--------+
                    | Network-1 | | Network-2 +---- fd00:2::/48
                    +------+----+ +-----+-----+
                            \          /
                             \        /
              2001:db8:1::/64 \      / 2001:db8:2::/64
                           +---+----+---+
                           |    Host    |
                           +------------+
```

4.2.1. Processing Source Address Selection Policy

The two ISPs, Network-1 and Network-2, provide the source address selection policy below,

Network-1 "if dst ::/0, then src 2001:db8:1::/64"
Network-2 "if dst ::/0 or fd00:2::/48, then src 2001:db8:2::/64"

and the Host's routing table looks like below,

Prefix	Nexthop
::/0	Network-1
fd00:2::/48	Network-2

As mentioned in the previous section, the policy that came from the selected entity in the routing table should be selected in the policy table also. In this case, the routing table selected Network-1 for the nexthop for the prefix ::/0, so the Network-1's policy should be selected in the policy table.

Prefix	Precedence	Label
::/0	?	1
2001:db8:1::/64	?	1
fd00:2::/48	?	5
2001:db8:2::/64	?	5

Next comes the merging with the default policy table, which is pasted below from the RFC 3484,

Prefix	Precedence	Label
::1/128	50	0
::/0	40	1
2002::/16	30	2
::/96	20	3
::ffff:0:0/96	10	4

We can have the following merged table.

	Prefix	Precedence	Label
	::1/128	50	0
	::/0	40	1
*S	2001:db8:1::/64	?	1
*S	fd00:2::/48	?	5
*S	2001:db8:2::/64	?	5
	2002::/16	30	2
	::/96	20	3
	::ffff:0:0/96	10	4

There are unresolved values in the table, which should be resolved in accordance with other policies below.

4.2.2. Processing Destination Address Selection Policy

The above mentioned two ISPs, Network-1 and Network-2, also want to provide the destination address selection policy below,

Network-1 "::/0 Precedence 20, ::ffff:0:0/96 Precedence 10"
Network-2 "::/0 Precedence 30, ::ffff:0:0/96 Precedence 40"

and the routing table of the Host is like below.

Prefix	Nexthop
::/0	Network-1
::ffff:0:0/96	Network-2
fd00:2::/48	Network-2

Here, the merging process selects the Precedence value of the policy that is selected in the routing table. That is, the routing table above selects Network-1 for the prefix ::/0, so the Precedence value for the prefix ::/0 should be the value of the Network-1's policy. The resulf of this merging process is below.

Prefix	Precedence	Label
::/0	20	?
::ffff:0:0/96	40	?

Next comes the merging with the policy table that is merged with the source address selection policy.

	Prefix	Precedence	Label
	::1/128	50	0
*D	::/0	20	1
*S	2001:db8:1::/64	?	1
*S	fd00:2::/48	?	5
*S	2001:db8:2::/64	?	5
	2002::/16	30	2
	::/96	20	3
*D	::ffff:0:0/96	40	4

4.2.3. Precedence and Label Values Adjustment

Regarding the unresolved value in the previous unfinished policy table, the merging process should choose the most harmless Precedence value. That means, the Precedence value that does not spoil or change the other policy table entries' effects.

The process should find a prefix that best includes or matches the prefix with the unresolved value in this policy table, and use the Precedence value of the selected prefix. In this example, the prefix 2001:db8:1::/64 longestly matches with the existing prefix ::/0, so the Precedence value 20 was used for the merged entries. The same value goes to the entries with 2001:db8:2::/64 and fd00:2::/48.

	Prefix	Precedence	Label
	::1/128	50	0
*D	::/0	20	1
*S	2001:db8:1::/64	20	1
*S	fd00:2::/48	20	5
*S	2001:db8:2::/64	20	5
	2002::/16	30	2
	::/96	20	3
*D	::ffff:0:0/96	40	4

Though not ducumented in this example, almost the same process as the Precedence value adjustment should be applied to the Label value. That is, the merging process should choose the most harmless Label value, which does not spoil or change the other policy table entries' effects. Hense, it should use the same Label value as the existing entry that longestly matches the prefix of the unresolved Label value.

5. Discussion

In this document, we examined and classified address selection policies. For each kind, we proposed how to solve the merging conflicting policies.

As documented here, the merging process has close relationship with the routing table. It should be noted that the address selection policy distribution has to be considered and used along with the routing mechanism.

6. IANA Considerations

This document has no actions for IANA.

7. Security Considerations

TBD

8. Acknowledgements

Dave Thaler and Aleksi Suhonen has given invaluable advice and feedback on this document.

9. References

9.1. Normative References

[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.

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

[RFC5220] Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Problem Statement for Default Address Selection in Multi-Prefix Environments: Operational Issues of RFC 3484 Default Rules", RFC 5220, July 2008.

9.2. Informative References

[I-D.ietf-6man-addr-select-considerations]
Chown, T., "Considerations for IPv6 Address Selection Policy Changes", draft-ietf-6man-addr-select-considerations-00 (work in progress), October 2009.

Appendix A. Revision History

02:

The document structure was changed.
Detailed the whole picture of merging process implementation in Section 4.
Documented how the source address policy merging and destination address policy merging interacts in Section 4.

01:

The section 4 was made clearer.
The section 5 "Conceptual processing model" was added.
The section "Conclusions" was renamed to "Discussion".

Authors' Addresses

Arifumi Matsumoto
NTT PF Lab
Midori-Cho 3-9-11
Musashino-shi, Tokyo 180-8585
Japan

Phone: +81 422 59 3334
Email: arifumi@nttv6.net

Tomohiro Fujisaki
NTT PF Lab
Midori-Cho 3-9-11
Musashino-shi, Tokyo 180-8585
Japan

Phone: +81 422 59 7351
Email: fujisaki@syce.net

Ruri Hiromi
Intec Netcore, Inc.
Shinsuna 1-3-3
Koto-ku, Tokyo 136-0075
Japan

Phone: +81 3 5665 5069
Email: hiromi@inetcore.com