IPv6 Operations Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 14, 2007

A. Matsumoto
T. Fujisaki
NTT
R. Hiromi
K. Kanayama
Intec Netcore
November 10, 2006

# Problem Statement of Default Address Selection in Multi-prefix Environment: Operational Issues of RFC3484 Default Rules
## draft-ietf-v6ops-addr-select-ps-00.txt



Abstract

One physical network can carry multiple logical networks. Moreover, we can use multiple physical networks at the same time in a host. In that environment, end-hosts might have multiple IP addresses and be required to use them selectively. Without an appropriate source/

destination address selection mechanism, the host will experience some trouble in the communication. RFC 3484 defines both the source and destination address selection algorithms, but the multi-prefix environment considered here needs additional rules beyond the default operation. This document describes the possible problems that end-hosts could encounter in an environment with multiple logical networks.

Table of Contents

## 1. Introduction

One physical network can carry multiple logical networks. In that case, an end-host has multiple IP addresses. In the IPv4-IPv6 dual stack environment or in a site connected to both ULA [RFC4193] and Global scope networks, an end-host has multiple IP addresses. These are examples of the networks that we focus on in this document. In such an environment, an end-host will encounter some communication trouble.

Inappropriate source address selection at the end-host causes unexpected asymmetric routing or filtering by a router on the way back or discard due to there being no route to the host.

Considering a multi-prefix environment, the destination address selection is also important for correct communication establishment. The key to the appropriate process will come from the way to configure the source address and destination address to the interfaces at the end-hosts by the network policy of the site.

RFC 3484 [RFC3484] defines both source and destination address selection algorithms. In most cases, the host will be able to communicate with the targeted host using the algorithms. But there are still problematic cases such as when multiple default routes are supplied. This document describes such possibilities of false dropping during address selection.

In addition, the provision of an address policy table is an important matter. RFC 3484 describes all the algorithms for setting the address policy table but it makes no mention of the provisions of address policy and does not define how to set it except manually.

### 1.1. Scope of this document

There has been a lot of discussion about "multiple addresses/ prefixes" but the multi-homing issues for redundancy are out of our scope. Cooperation with a mechanism like shim6 is rather desirable. We focus on an end-site network environment. The scope of this document is to sort out problematic cases of false dropping of the address selection within a multi-prefix environment.

## 2. Problem Statement

### 2.1. Source Address Selection

### 2.1.1. Multiple Routers on Single Interface

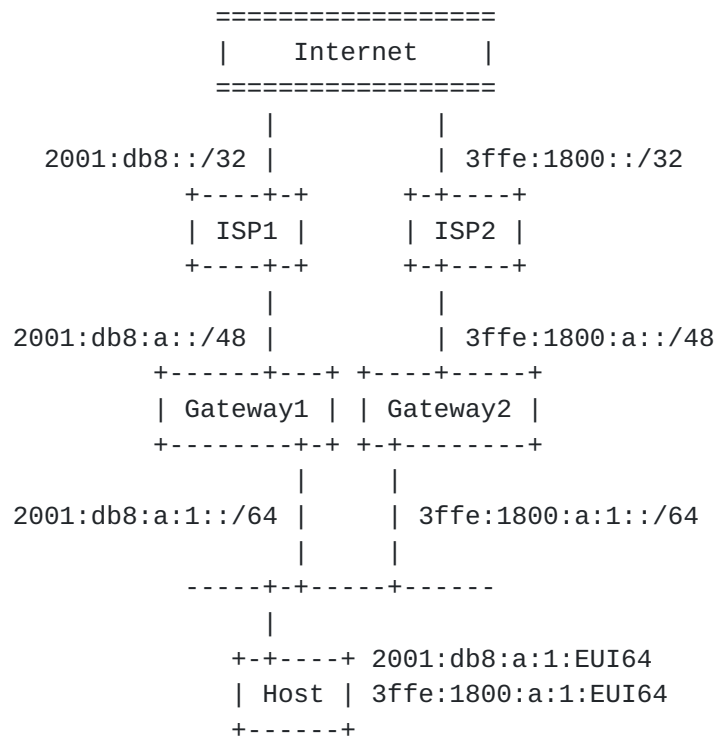


[Fig. 1]

Generally speaking, there is no interaction between next-hop determination and address selection. In this example, when Host sends a packet via Gateway1, the Host does not necessarily choose the address 2001:db8:a:1::EUI64 given by Gateway1 as the source address. This causes the same problem as described in the next section 'Ingress Filtering Problem'.

To solve this case, one approach is to configure correctly both the routing configuration and address selection policy at Host. You can use RFC 4191 [RFC4191] to deliver routing information to hosts. Another approach is to configure the gateways to make use of packet redirection between the gateways.

### 2.1.2. Ingress Filtering Problem

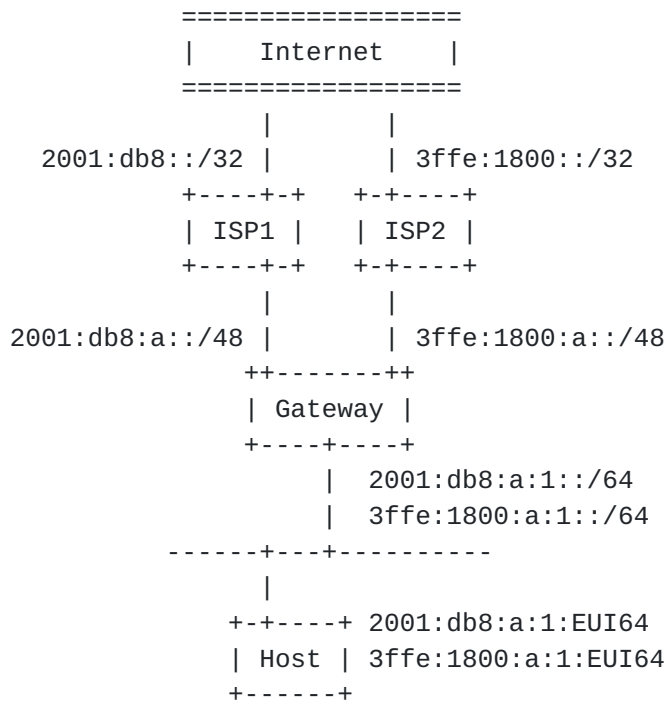


[Fig. 2]

When a relatively small site, which we call a "customer network", is attached to two upstream ISPs, each ISP delegates a network address block, which is usually /48, and a host has multiple IPv6 addresses.

When the source address of an outgoing packet is not the one that is delegated by an upstream ISP, there is a possibility that the packet will be dropped at the ISP by its Ingress Filter. Ingress filtering(uRPF: unicast Reverse Path Forwarding) is becoming more and more popular among ISPs in order to mitigate the damage of DoS attacks.

In this example, when the Gateway chooses the default route to ISP2 and the Host chooses 2001:db8:a:1::EUI64 as the source address for packets sent to a host(2001:fa8::1) somewhere in the Internet, the packets may be dropped at ISP2 because of Ingress Filtering.

One possible solution for this problem is to adopt source-address-based routing at the customer site's gateway, but this manner of routing is not very popular at the moment.

### 2.1.3. Half-Closed Network Problem

You can see a second typical source address selection problem in a multihome site with global-closed mixed connectivity like the figure below. In this case, Host-A is in a multihomed network and has two IPv6 addresses, one delegated from each of the upstream ISPs. Note that ISP2 is a closed network and does not have connectivity to the Internet.

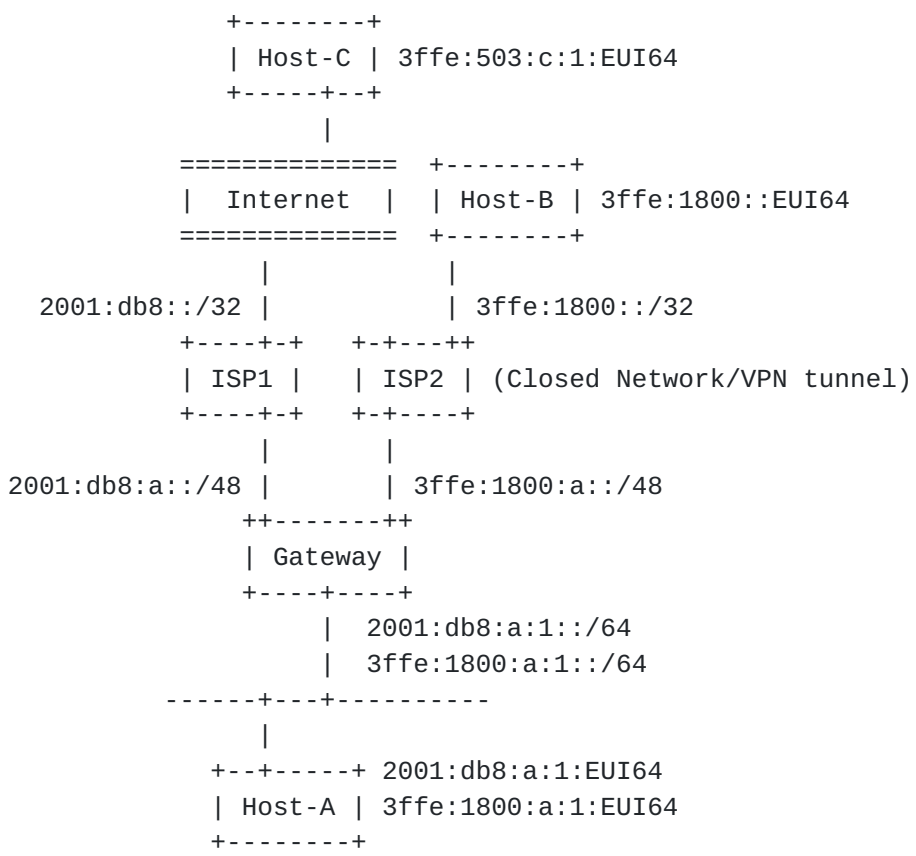


[Fig. 3]

You don't need two physical network connection here. The connection from Gateway to ISP2 can be a logical link over ISP1 and the Internet.

When Host-A starts the connection to Host-B in ISP2, the source address of a sending packet will be the one delegated from ISP2, that is 3ffe:1800:a:1:EUI64, because of rule 8 (longest matching prefix) in RFC 3484.

Host-C is located somewhere in the Internet and has an IPv6 address 3ffe:503:c:1:EUI64. When Host-A sends a packet to Host-C, the longest matching algorithm chooses 3ffe:1800:a:1:EUI64 for the source

address. In this case, the packet goes through ISP1 and may be filtered by ISP1's ingress filter. Even if the packet is fortunately not filtered by ISP1, a return packet from Host-C cannot possibly be delivered to Host-A because the return packet is destined for 3ffe:1800:a:1:EUI64, which is closed from the Internet.

In this case, source-address-based routing alone described in the previous section does not solve the problem. What is important is that each host chooses a correct source address for a given destination address as far as NAT does not exist in the IPv6 world.

### 2.1.4. Combined Use of Global and ULA

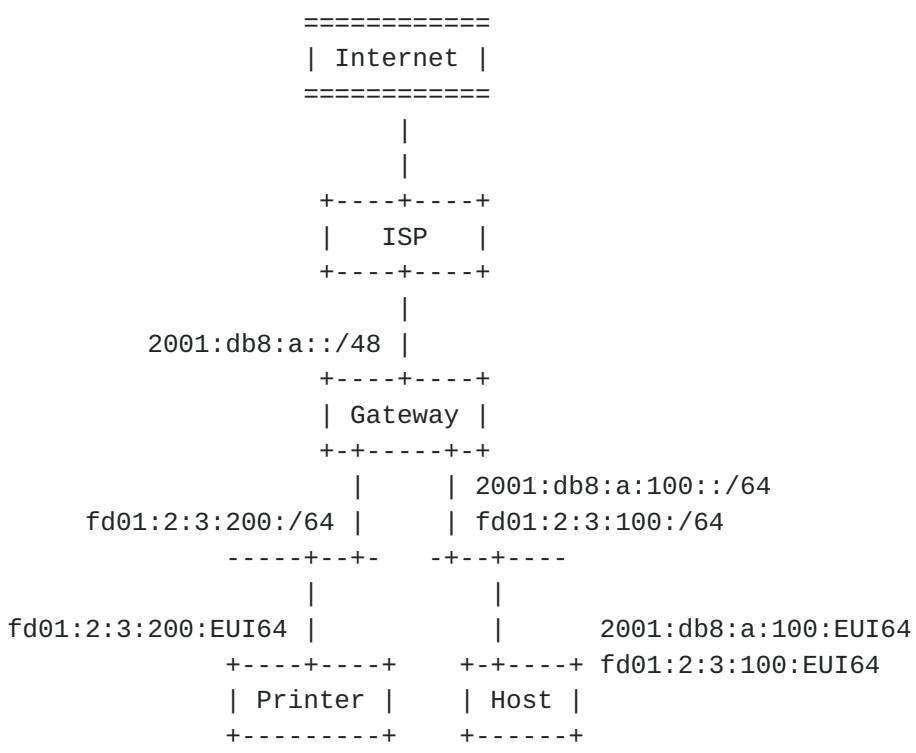


[Fig. 4]

As NAP [I-D.ietf-v6ops-nap] describes, using ULA may be beneficial in some scenarios. If ULA is used for internal communication, packets with ULA addresses need to be filtered at Gateway.

There is no serious problem related to address selection in this case, thanks to the unlikeness of ULA and Global Unicast Address for now. RFC 3484's longest matching rule chooses the correct address for both intra-site and extra-site communication.

In a few years, however, the longest matching rule will not be able to choose the correct address anymore: the moment the assignment of

those Global Unicast Addresses whose beginning bit is 1 starts. In RFC 4291 [RFC4291], almost all the space of IPv6, including those with beginning bit 1, is assigned as Global Unicast Addresses.

### 2.1.5. Site Renumbering

RFC 4192 [RFC4192] describes a recommended procedure for renumbering a network from one prefix to another. An auto-configured address has a lifetime, so by stopping advertisement of the old prefix it is eventually invalidated.

However, it takes a long time to invalidate the old prefix. You cannot stop routing to the old prefix as long as the old prefix is not deprecated. This issue can be a tough issue for ISP network administrator.

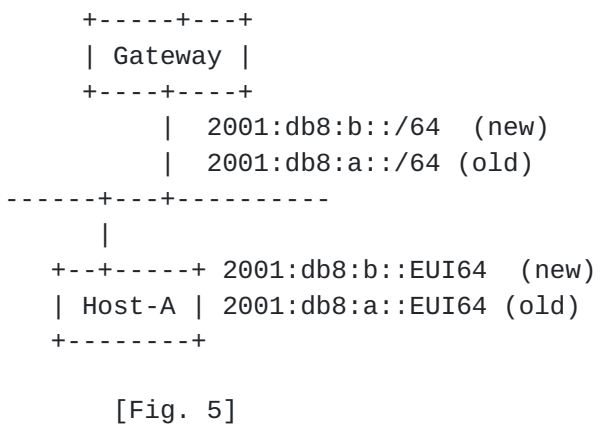


[Fig. 5]

### 2.1.6. Multicast Source Address Selection

This case is an example of Site-local or Global prioritization. When you send a multicast packet across site-borders, the source address of the multicast packet must be a global scope address. The longest matching algorithm, however, selects a ULA address if the sending host has both a ULA and a global address.

### 2.1.7. Temporary Address Selection

RFC 3041 [RFC3041] defines a Temporary Address. The usage of Temporary Address has both pros and cons. It is good for viewing web-pages or communicating with the general public, but it is bad for a service that uses address-based authentication and for logging purpose.

It would be better if you could turn the temporary address on and off. It would also be better if you could switch its usage per service(destination address). The same situation can be found when

using HA and CoA in MobileIP network.

## 2.2. Destination Address Selection

### 2.2.1. IPv4 or IPv6 prioritization

The default policy table gives IPv6 addresses higher precedence than IPv4 addresses. There seem to be many cases, however, where network administrators want to control the address selection policy of end-hosts the other way around.

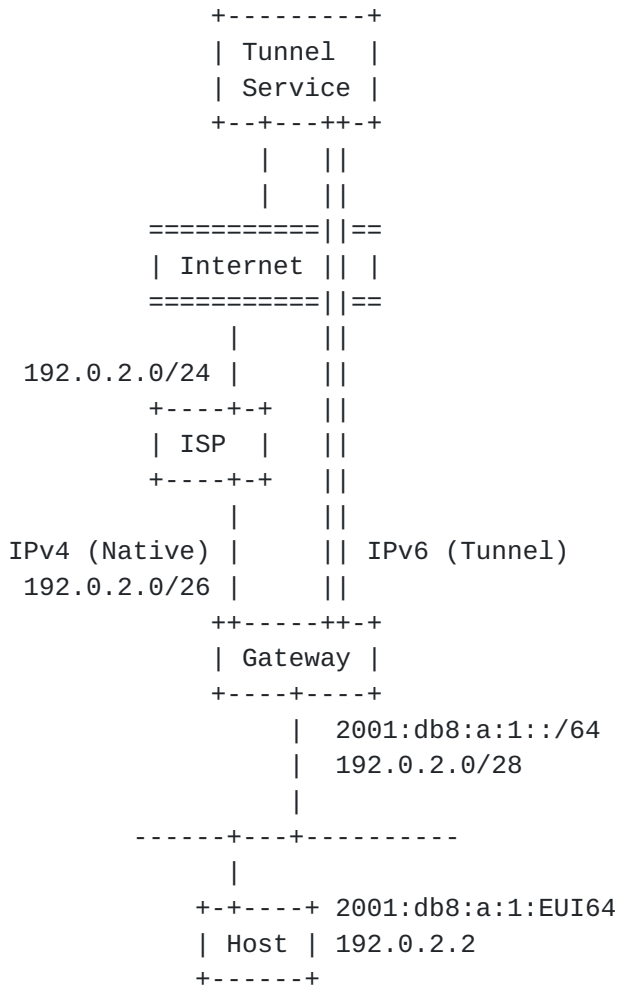


[Fig. 6]

In the figure above, a site has native IPv4 and tunneled IPv6 connectivity. Therefore, the administrator may want to set a higher priority for using IPv4 than using IPv6 because the quality of the tunnel network seems to be worse than that of the native transport.

### 2.2.2. ULA and IPv4 dual-stack environment

This is a special form of IPv4 and IPv6 prioritization. When an enterprise has IPv4 Internet connectivity but does not yet have IPv6 Internet connectivity, and the enterprise wants to provide site-local IPv6 connectivity, ULA is the best choice for site-local IPv6 connectivity. Each employee host will have both an IPv4 global or private address and a ULA. Here, when this host tries to connect to Host-C that has registered both A and AAAA records in the DNS, the host will choose AAAA as the destination address and ULA for the source address. This will clearly result in a connection failure.

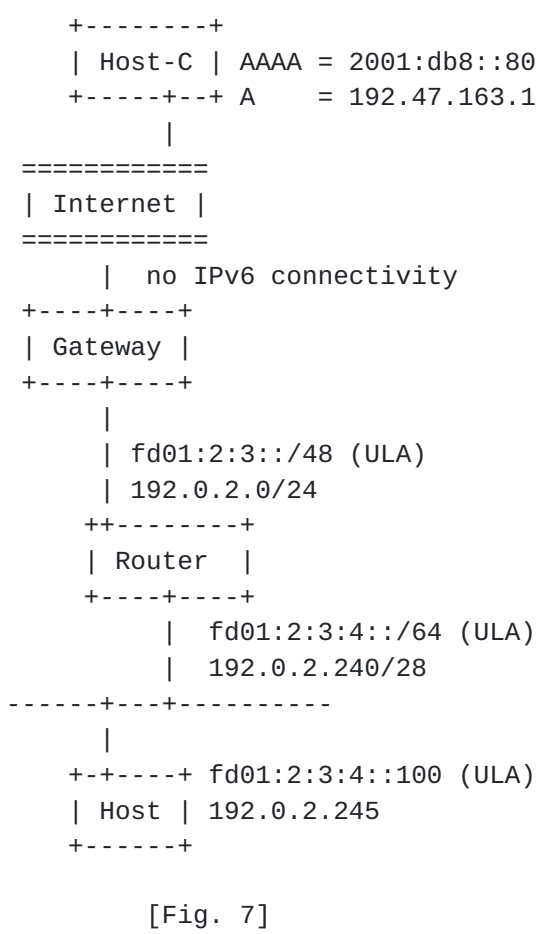


[Fig. 7]

### 2.2.3. ULA or Global Prioritization

It is very common to differentiate services by the client's source address. IP-address-based authentication is an extreme example of this. Another typical example is a web service that has pages for the public and internal pages for employees or involved parties. Yet another example is DNS zone splitting.

However, ULA and IPv6 global address both have global scope, and RFC

3484 default rules do not specify which address should be given priority. This point makes IPv6 implementation of address-based service differentiation a bit harder.

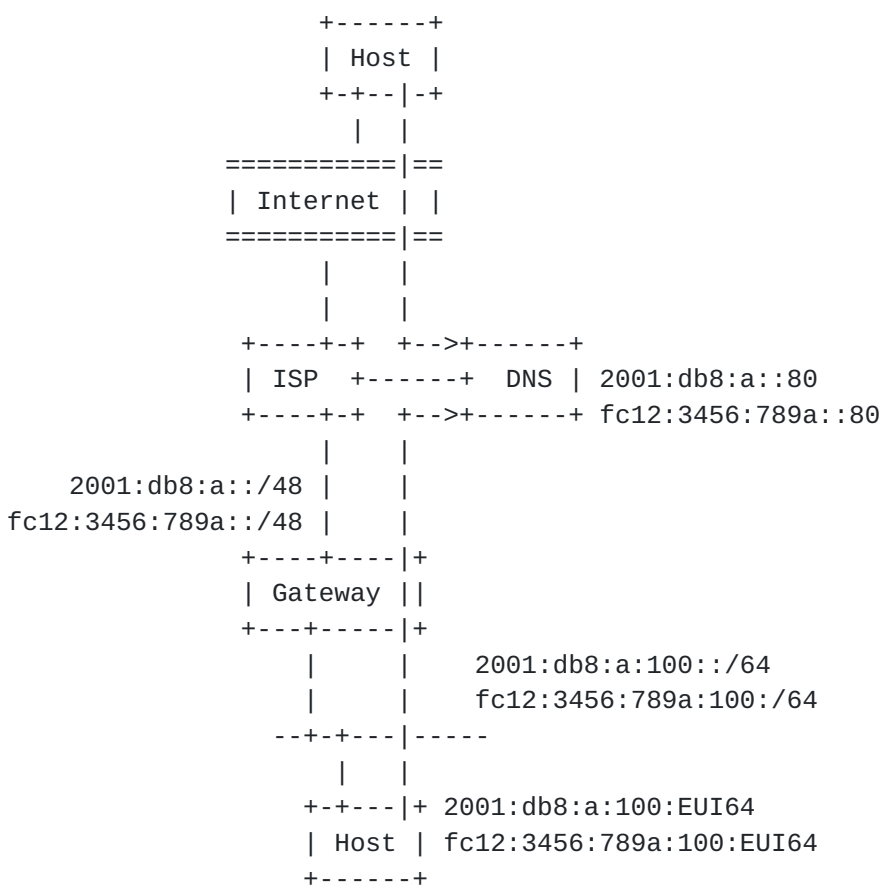


[Fig. 7]

## 3. Solutions

### 3.1. More Specific Routes (RFC 4191)

This method enables network administrator to distribute routing information to end-hosts. It can solve only two problems in this document, that is 2.1.1, 2.2.2. Routing information doesn't determine the source address when multiple addresses are attached to the outgoing network interface. So, it cannot be used for every cases here.

### 3.2. Policy Table Manipulation

Almost all the problem cases raised in this document can be solved by configuring the policy table at end-hosts. The problem for a site-

administrator is that he does not have the means to deliver policies to end-hosts. Therefore, we proposed a method for policy distribution in the form of DHCPv6 option [I-D.fujisaki-dhc-addr-select-opt]. The usage of this mechanisim is illustrated in another I-D [I-D.arifumi-ipv6-policy-dist].

### 3.3. Revising RFC 3484

Revising address selection rules defined in RFC 3484 in another idea. These problems are, however, too network-environment-specific, so it's not easy to have all-purpose rule set.

## 4. Conclusion

We have covered problems related to destination or source address selection. These problems have their roots in the situation where end-hosts have multiple IP addresses. In this situation, every end-host must choose an appropriate destination and source address, which cannot be achieved only by routers.

It should be noted that end-hosts must be informed about routing policies of their upstream networks for appropriate address selection. A site administrator must consider every possible address false-selection problem and take countermeasures beforehand.

## 5. Security Considerations

Address false-selection can lead to serious security problem, such as session hijack. However, it should be noted that address selection is eventually up to end-hosts. We have no means to enforce one specific address selection policy to every end-host. So, a network administrator has to take countermeasures for unexpected address selection.

## 6. IANA Considerations

This document has no actions for IANA.

## 7. References

### 7.1. Normative References

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

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.

## 7.2. Informative References

[I-D.arifumi-ipv6-policy-dist]
Matsumoto, A., "Practical Usages of Address Selection Policy Distribution", draft-arifumi-ipv6-policy-dist-01 (work in progress), June 2006.

[I-D.fujisaki-dhc-addr-select-opt]
Fujisaki, T., "Distributing Default Address Selection Policy using DHCPv6", draft-fujisaki-dhc-addr-select-opt-02 (work in progress), June 2006.

[I-D.ietf-v6ops-nap]
Velde, G., "Network Architecture Protection for IPv6", draft-ietf-v6ops-nap-04 (work in progress), October 2006.

[RFC3041] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.

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

[RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, February 2006.

## Appendix A. Appendix. Revision History

01:
Authors' addresses corrected.
Solutions section added.
Security Considerations section fully rewritten.
Some editorial changes.

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

Ken-ichi Kanayama
Intec Netcore, Inc.
Shinsuna 1-3-3
Koto-ku, Tokyo 136-0075
Japan

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

## Full Copyright Statement



## Intellectual Property



## Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).