

6man Working Group
Internet-Draft
Intended status: Standards Track
Expires: July 20, 2013

A. Matsumoto
T. Fujisaki
NTT
T. Chown
University of Southampton
January 16, 2013

Distributing Address Selection Policy using DHCPv6
draft-ietf-6man-addr-select-opt-08.txt

Abstract

[RFC 6724](#) defines default address selection mechanisms for IPv6 that allow nodes to select an appropriate address when faced with multiple source and/or destination addresses to choose between. The [RFC 6724](#) allowed for the future definition of methods to administratively configure the address selection policy information. This document defines a new DHCPv6 option for such configuration, allowing a site administrator to distribute address selection policy overriding the default address selection parameters and policy table, and thus control the address selection behavior of nodes in their site.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 20, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

1. Introduction

[RFC 3484](#) [[RFC3484](#)] describes default algorithms for selecting an address when a node has multiple destination and/or source addresses to choose from by using an address selection policy. In [Section 2 of RFC 6724](#), it is suggested that the default policy table may be administratively configured to suit the specific needs of a site. This specification defines a new DHCPv6 option for such configuration.

Some problems have been identified with the default [RFC 3484](#) address selection policy [[RFC5220](#)]. It is unlikely that any default policy will suit all scenarios, and thus mechanisms to control the source address selection policy will be necessary. Requirements for those mechanisms are described in [[RFC5221](#)], while solutions are discussed in [[I-D.ietf-6man-addr-select-sol](#)] and [[I-D.ietf-6man-addr-select-considerations](#)]. Those documents have helped shape the improvements in the default address selection algorithm [[RFC6724](#)] as well as the DHCPv6 option defined in this specification.

1.1. Conventions Used in This Document

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](#)].

1.2. Terminology

This document uses the terminology defined in [\[RFC2460\]](#) and the DHCPv6 specification defined in [\[RFC3315\]](#)

2. Address Selection options

The Address Selection option provides the address selection policy table, and some other configuration parameters.

An Address Selection option contains zero or more policy table options. Multiple Policy Table options in an Address Selection option constitute a single policy table.

The format of the Address Selection option is given below.




Figure 1: Address Selection option format

option-code: OPTION_ADDRSEL (TBD).

option-len: The total length of the Reserved field, A, P flags, and POLICY TABLE OPTIONS in octets.

Reserved: Reserved field. Server MUST set this value to zero and client MUST ignore its content.

A: Automatic Row Addition flag. This flag toggles the Automatic Row Addition flag at client hosts, which is described in the [section 2.1 in RFC 6724](#) [\[RFC6724\]](#). If this flag is set to 1, it does not change client host behavior, that is, a client MAY automatically add additional site-specific rows to the policy table. If set to 0, the Automatic Row Addition flag is disabled, and a client SHOULD NOT automatically add rows to the policy table.

P: Privacy Preference flag. This flag toggles the Privacy Preference flag at client hosts, which is described in the [section 5 in RFC 6724](#) [RFC6724]. If this flag is set to 1, it does not change client host behavior, that is, a client will prefer temporary addresses. If set to 0, the Privacy Preference flag is disabled, and a client will prefer public addresses.

POLICY TABLE OPTIONS: Zero or more Address Selection Policy Table options described below.




Figure 2: Address Selection Policy Table option format

option-code: OPTION_ADDRSEL_TABLE (TBD).

option-len: The total length of the label field, precedence field, prefix-len field, and prefix field.

label: An 8-bit unsigned integer; this value is for correlation of source address prefixes and destination address prefixes.

precedence: An 8-bit unsigned integer; this value is used for sorting destination addresses.

prefix-len: An 8-bit unsigned integer; the number of leading bits in the prefix that are valid. The value ranges from 0 to 128.

prefix: A variable-length field containing an IP address or the prefix of an IP address. An IPv4-mapped address [RFC4291] must be used to represent an IPv4 address as a prefix value. The prefix should be left aligned, big-endian, and zero padded on the right up to the next octet boundary. So the length of this field should be between 0 and 16 bytes.

3. Appearance of the Address Selection options

The Address Selection options MUST NOT appear in any DHCPv6 messages other than the following ones: Solicit, Advertise, Request, Renew, Rebind, Reconfigure, Information-Request, and Reply.

4. Processing the Policy Table option

This section describes how to process received Policy Table option at the DHCPv6 client.

This option's concept is to serve as a hint for a node about how to behave in the network. So, basically, it should be up to the node's administrator how to deal with the received policy information in the way described below.

4.1. Handling of the local policy table

[RFC 6724](#) defines the default policy table. Also, users are usually able to configure the policy table to satisfy their own requirements.

The client implementation SHOULD provide the following choices to the user:

- a) replace the existing active policy table with the DHCPv6 distributed policy table.
- b) preserve the existing active policy table, whether this be the default policy table, or user configured policy.

4.2. Handling of the stale policy table

When the information from the DHCP server goes stale, the policy received from the DHCP server should be deprecated.

The received information can be considered stale in several cases, such as, when the interface goes down, the DHCP server does not respond for a certain amount of time, and the Information Refresh Time is expired.

4.3. Multi-interface situation

The policy table, and other parameters specified in this document are node-global information by their nature. One reason being that the outbound interface is usually chosen after destination address selection. So, a host cannot make use of multiple address selection policies even if they are stored per interface.

Even if the received policy from one source is merged with one from another source, the effect of both policy are more or less changed. The policy table is defined as a whole, so the slightest addition/deletion from the policy table brings a change in semantics of the policy.

It also should be noted that absence of the distributed policy from a certain network interface should not be treated as absence of policy itself, because it may mean preference for the default address selection policy.

Under the above assumptions, how to handle received policy is specified below.

A node MAY use Address Selection options by default in any of the following two cases:

- 1: The host is single-homed, where the host belongs to one administrative network domain exclusively usually through one active network interface.
- 2: The host implements some advanced heuristics to deal with multiple received policy, which is outside the scope of this document.

The above restrictions do not preclude implementations from providing configuration options to enable this option on a certain network interface.

Nor, they do not preclude implementations from storing distributed address selection policies per interface. They can be used effectively on such implementations that adopt per-application interface selection.

5. Implementation Considerations

- o The value 'label' is passed as an unsigned integer, but there is no special meaning for the value, that is whether it is a large or small number. It is used to select a preferred source address prefix corresponding to a destination address prefix by matching the same label value within the DHCP message. DHCPv6 clients SHOULD convert this label to a representation appropriate for the local implementation (e.g., string).
- o Currently, the label and precedence values are defined as 8-bit unsigned integers. In almost all cases, this value will be enough.

- o The maximum number of address selection rules that may be conveyed in one DHCPv6 message depends on the prefix length of each rule and the maximum DHCPv6 message size defined in [RFC 3315](#). It is possible to carry over 3,000 rules in one DHCPv6 message (maximum UDP message size). However, it should not be expected that DHCP clients, servers and relay agents can handle UDP fragmentation. Network administrators SHOULD consider local limitations to the maximum DHCPv6 message size that can be reliably transported via their specific local infrastructure to end nodes; and therefore they SHOULD consider the number of options, the total size of the options, and the resulting DHCPv6 message size, when defining their Policy Table.
- o Since the number of selection rules could be large, an administrator configuring the policy to be distributed should consider the resulting DHCPv6 message size.

6. Security Considerations

A rogue DHCPv6 server could issue bogus address selection policies to a client. This might lead to incorrect address selection by the client, and the affected packets might be blocked at an outgoing ISP because of ingress filtering, incur additional network charges, or be misdirected to an attacker's machine. Alternatively, an IPv6 transition mechanism might be preferred over native IPv6, even if it is available. To guard against such attacks, a legitimate DHCPv6 server should communicate through a secure, trusted channel, such as a channel protected by IPsec, SEND and DHCP authentication, as described in [section 21 of RFC 3315](#),

Another threat is about privacy concern. As in the security consideration section of [RFC 6724](#), at least a part of, the address selection policy stored in a host can be leaked by a packet from a remote host. This issue will not be modified by the introduction of this option, regardless of whether the host is multihomed or not.

7. IANA Considerations

IANA is requested to assign option codes to OPTION_ADDRSEL and OPTION_ADDRSEL_TABLE from the option-code space as defined in section "DHCPv6 Options" of [RFC 3315](#).

8. References

8.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", [RFC 3315](#), July 2003.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", [RFC 3484](#), February 2003.
- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", [RFC 6724](#), September 2012.

8.2. Informative References

- [I-D.ietf-6man-addr-select-considerations]
Chown, T. and A. Matsumoto, "Considerations for IPv6 Address Selection Policy Changes", [draft-ietf-6man-addr-select-considerations-04](#) (work in progress), October 2011.
- [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.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC3493] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W. Stevens, "Basic Socket Interface Extensions for IPv6", [RFC 3493](#), February 2003.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.
- [RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", [RFC 4941](#), September 2007.
- [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.

[RFC5221] Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Requirements for Address Selection Mechanisms", [RFC 5221](#), July 2008.

[Appendix A.](#) Acknowledgements

Authors would like to thank to Dave Thaler, Pekka Savola, Remi Denis-Courmont, Francois-Xavier Le Bail, Ole Troan, Bob Hinden, Dmitry Anipko, Ray Hunter, Rui Paulo, Brian E Carpenter, Tom Petch, and the members of 6man's address selection design team for their invaluable contributions to this document.

[Appendix B.](#) Past Discussion

- o The 'zone index' value is used to specify a particular zone for scoped addresses. This can be used effectively to control address selection in the site scope (e.g., to tell a node to use a specified source address corresponding to a site-scoped multicast address). However, in some cases such as a link-local scope address, the value specifying one zone is only meaningful locally within that node. There might be some cases where the administrator knows which clients are on the network and wants specific interfaces to be used though. However, in general case, it is really rare case, and the field was removed.

Authors' Addresses

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

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

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

Phone: +81 422 59 7351
Email: fujisaki@nttv6.net

Tim Chown
University of Southampton
Southampton, Hampshire SO17 1BJ
United Kingdom

Email: tjc@ecs.soton.ac.uk

