DHC Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 9, 2010

T. Fujisaki
A. Matsumoto
NTT
R. Hiromi
Intec Netcore
March 8, 2010

Distributing Address Selection Policy using DHCPv6
draft-fujisaki-dhc-addr-select-opt-09.txt

Abstract

This document describes a new DHCPv6 option for distributing address selection policy information defined in RFC3484 and/or source address and destination address selection policies to a client. With this option, site administrators can distribute address selection policy to control the node's address selection behavior.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on September 9, 2010.

Copyright Notice

Copyright (c) 2010 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 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

RFC3484 [RFC3484] describes algorithms for selecting a default address when a node has multiple destination and/or source addresses by using an address selection policy. However, there are some problems with the default address selection policy in RFC3484 [RFC5220], and mechanisms to control a proper source address selection will be necessary. Requirements for those mechanisms are described in [RFC5221], and solutions are discussed in [I-D.ietf-6man-addr-select-sol]. This document describes an option for distributing address selection policy information using DHCPv6, which is referred as `most proactive approach' in the solution document.

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 [RFC2119].

1.2. Terminology

This document uses the terminology defined in [RFC2460] and the DHCP specification defined in [RFC3315]

2. Address Selection Policy Option

The Address Selection Policy Option provides policy information for address selection rules. Specifically, it transmits a set of IPv6 source and destination address prefixes and some parameters that are used to control address selection as described in RFC 3484.

Each end node is expected to configure its policy table, as described in RFC 3484, using the Address Selection Policy option information as an reference.

The format of the Address Selection Policy option is given below:

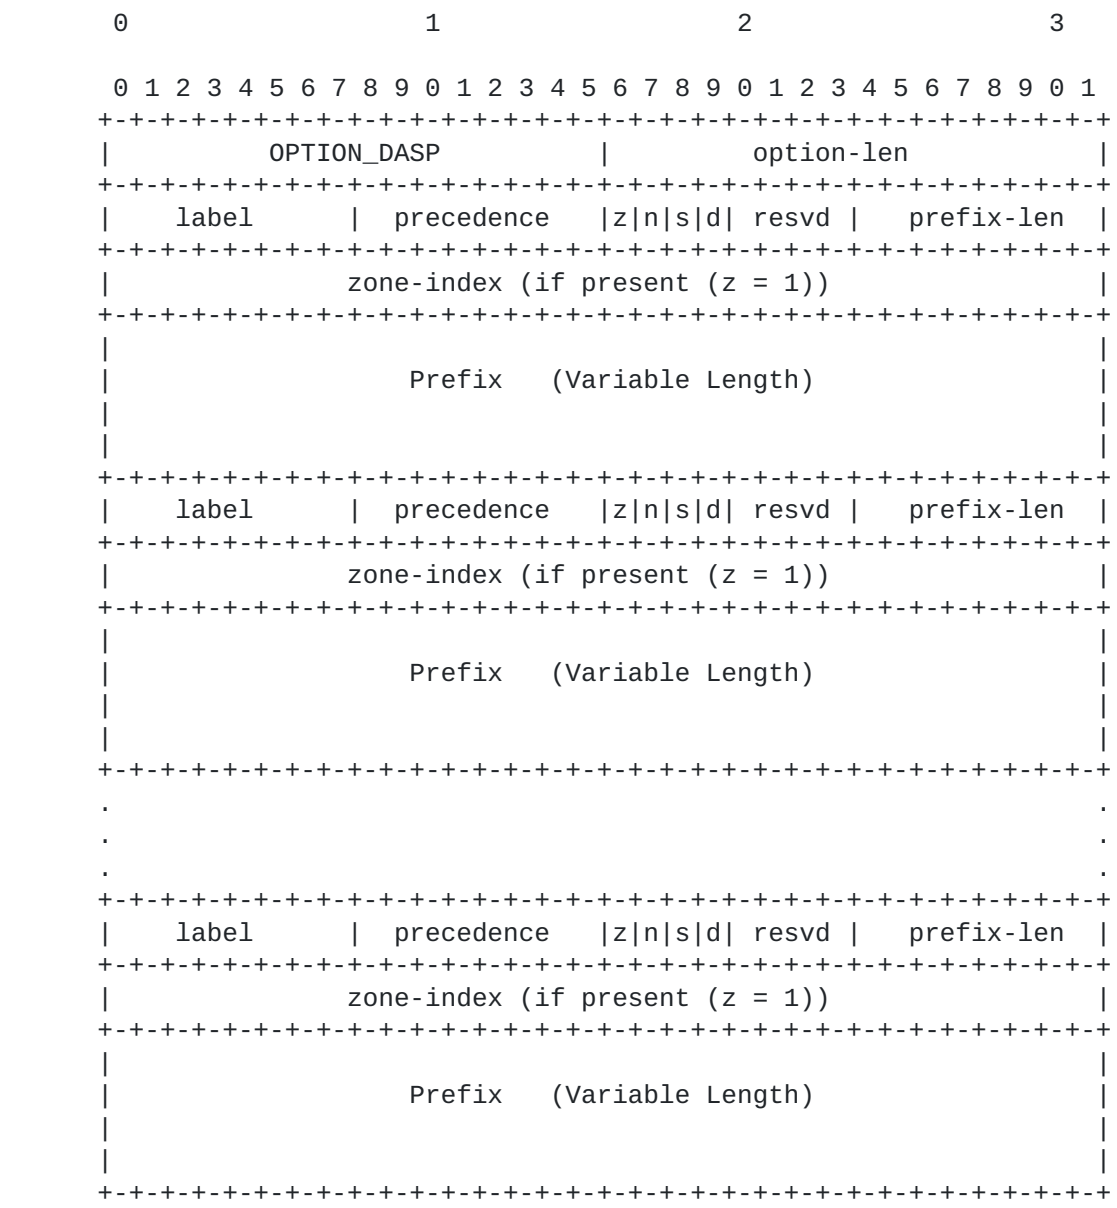
0 1 2 3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
OPTION_DASP
option-len
label
precedence
z|n|s|d
resvd
prefix-len
zone-index (if present (z = 1))
Prefix (Variable Length)
label
precedence
z|n|s|d
resvd
prefix-len
zone-index (if present (z = 1))
Prefix (Variable Length)
label
precedence
z|n|s|d
resvd
prefix-len
zone-index (if present (z = 1))
Prefix (Variable Length)

[Fig. 1]

Fields:

option-code: OPTION_DASP (TBD)

option-len: The total length of the label fields, precedence fields, zone-index fields, prefix-len fields, and prefix fields in octets.

label: An 8-bit unsigned integer; this value is used to make a combination of source address prefixes and destination address prefixes.

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

z bit 'zone-index' bit. If z bit is set to 1, 32 bit zone-index value is included right after the "prefix-len" field, and "Prefix" value continues after the "zone-index" field. If z bit is 0, "Prefix" value continues right after the "prefix-len" value.

n bit 'no privacy iid' bit. If n bit is set to 1, RFC 4941 [RFC4941] privacy extensions MUST not be used for this prefix. If n bit is 0, interface ID may use RFC4941.

s bit 'source address selection policy' bit. If s bit is set to 1, this prefix is source address selection policy, not RFC3484 policy table entry. The usage of this policy is defined in [I-D.arifumi-6man-addr-select-conflict].

d bit 'destination address selection policy' bit. If d bit is set to 1, this prefix is destination address selection policy, not RFC3484 policy table entry. The usage of this policy is defined in [I-D.arifumi-6man-addr-select-conflict].

resvd 4-bit reserved field. Initialized to zero by sender, and ignored by receiver.

zone-index: If z-bit is set to 1, this field is inserted between "prefix-len" field and "Prefix" field. Zone-index field is an 32-bit unsigned integer and used to specify zones for scoped addresses. This bit length is defined in RFC3493 [RFC3493] as 'scope ID'.

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. The Prefix field is 0, 4, 8, 12, or 16 octets, depending on the length.

Prefix: A variable-length field containing an IP address or the prefix of an IP address. IPv4-mapped address [mapped] must be used to represent an IPv4 address as a prefix value.

3. Appearance of this Option

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

4. Implementation Considerations

If there are multiple DHCPv6 servers, a node may have multiple address selection policies. Since RFC3484 policy table is one and global for a node, multiple polices should be merged in one. In a case that node's interfaces belong to different management domain (e.g. each interfaces are connected different site), it would have conflict policies. The policy merging algorithm is defined in [I-D.arifumi-6man-addr-select-conflict].

- 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 this DHCP message. DHCPv6 clients need to convert this label to a representation specified by each implementation (e.g., string).

- Currently, the value label, precedence are defined as 8-bit unsigned integers. In almost all cases, this value will be enough.

- The 'precedence' is used to sort destination addresses. There might be some cases where precedence values will conflict when a client already has a selection policy configured or a client receives multiple policies from multiple DHCP servers (e.g., when a home gateway in a user network is connected to multiple upstream ISPs). In such cases, manual configuration of the policy will be necessary.

- The maximum number of address selection rules in one DHCPv6 message depend on the prefix length of each rules and maximum DHCPv6 message size defined in RFC3315. It is possible to carry over 3,000 rules (e.g. default policy table defined in RFC3484

contains 5 rules) in one DHCPv6 message (maximum UDP message size).

- Since the number of selection rules would be large, policy distributer should be care about the DHCPv6 message size.

5. Discussion

- 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 hard to use this value.

- Since we got a comment that some implementations use 32-bit integers for zone index value, we extended the bit length of the 'zone index' field. However, as described above, there might be few cases to specify 'zone index' in policy distribution, we defined this field as optional, controlled by a flag.

- There may be some demands to control the use of special address types such as the temporary addresses described in RFC4941 [RFC4941], address assigned by DHCPv6 and so on. (e.g., informing not to use a temporary address when it communicate within the an organization's network). It is possible to indicate the type of addresses using reserved field value.

- We also proposed a policy distribution option using a Router Advertisement message defined in RFC4861 [RFC4861]. There was a discussion that using DHCPv6 was more suitable to distribute a selection policy, because such policy should be distributed under the site administrator's centralized control.

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.

To guard against such attacks, both DHCP clients and servers SHOULD use DHCP authentication, as described in section 21 of RFC 3315, "Authentication of DHCP messages."

7. IANA Considerations

IANA is requested to assign option codes to OPTION_DASP from the option-code space as defined in section "DHCPv6 Options" of RFC 3315.

Appendix A. RFC3484 implementation status

Today, many operating systems implement address selection mechanism defined in RFC3484. Many of them, however, implement the specification partially. We summarize current implementation status of RFC 3484 at http://www.nttv6.net/dass/.

Appendix B. Revision History

09:
Add 's' and 'd' option for policy merge.
correct some typo.

08:
Add reference for policy conflict discussion.
Update some references.

07:
Added the n bit and its description.

06:
Added the reason to extend zone index field in discussions section.
References updated.
Authors' e-mail addresses corrected.
Some editorial changes.

05:
Extended bit length of the zone-index field to 32-bits (thank you Jinmei-sanfor your comment), and changed packet format to reflect the extension.
Refrect Yoshifuji-san's comment to use this option information as an reference.

Modified the text controlling special address types.

04:
Added description about policy merge.
Modified the text controlling special address types.

03:
Discussion about DHCPv6 packet size and number of rules added.
Authors' e-mail addresses corrected.
Some editorial changes.

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.

8.2. Informative References

[I-D.arifumi-6man-addr-select-conflict]
Matsumoto, A., Fujisaki, T., and R. Hiromi, "Considerations of address selection policy conflicts", draft-arifumi-6man-addr-select-conflict-01 (work in progress), October 2009.

[I-D.ietf-6man-addr-select-sol]
Matsumoto, A., Fujisaki, T., Hiromi, R., and K. Kanayama, "Solution approaches for address-selection problems", draft-ietf-6man-addr-select-sol-02 (work in progress), July 2009.

[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.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman,

"Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

[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.

Authors' Addresses

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

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

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

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