

Internet Engineering Task Force (IETF)
Request for Comments: 7078
Category: Standards Track
ISSN: 2070-1721

A. Matsumoto
T. Fujisaki
NTT
T. Chown
University of Southampton
January 2014

Distributing Address Selection Policy Using DHCPv6

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. RFC 6724 allows 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 allowing the administrator to control the address selection behavior of nodes in their site.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7078>.

Copyright Notice

Copyright (c) 2014 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

[RFC6724] 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. This specification defines a new DHCPv6 option for configuring the default policy table.

Some problems were identified with the default address selection policy as originally defined in [RFC3484]. As a result, RFC 3484 was updated and obsoleted by [RFC6724]. While this update corrected a number of issues identified from operational experience, 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 [ADDR-SEL]. Those documents have helped shape the improvements in the default address selection algorithm in [RFC6724] as well as the requirements for the DHCPv6 option defined in this specification.

This option's concept is to serve as a hint for a node about how to behave in the network. Ultimately, while the node's administrator can control how to deal with the received policy information, the implementation SHOULD follow the method described below uniformly to ease troubleshooting and to reduce operational costs.

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. When an Address Selection option does not contain a policy table option, it may be used to just convey the A and P flags for Automatic Row Additions and Privacy Preference, respectively.

The format of the Address Selection option is given below.

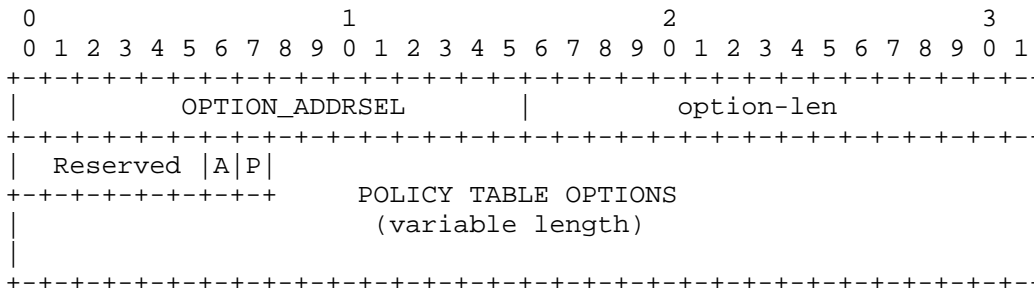


Figure 1: Address Selection Option Format

option-code: OPTION_ADDRSEL (84).

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

Reserved: Reserved field. The server MUST set this value to 0, and the 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 Section 2.1 of [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. If the option contains a POLICY TABLE option, this flag is meaningless,

and automatic row addition SHOULD NOT be performed against the distributed policy table. This flag SHOULD be set to 0 only when the Automatic Row Addition at client hosts is harmful for site-specific reasons.

P: Privacy Preference flag. This flag toggles the Privacy Preference flag on client hosts, which is described in Section 5 of [RFC6724]. If this flag is set to 1, it does not change client host behavior; that is, a client will prefer temporary addresses [RFC4941]. If set to 0, the Privacy Preference flag is disabled, and a client will prefer public addresses. This flag SHOULD be set to 0 only when the temporary addresses should not be preferred for site-specific reasons.

POLICY TABLE OPTIONS: Zero or more Address Selection Policy Table options, as described below. This option corresponds to a row in the policy table defined in Section 2.1 of [RFC6724].

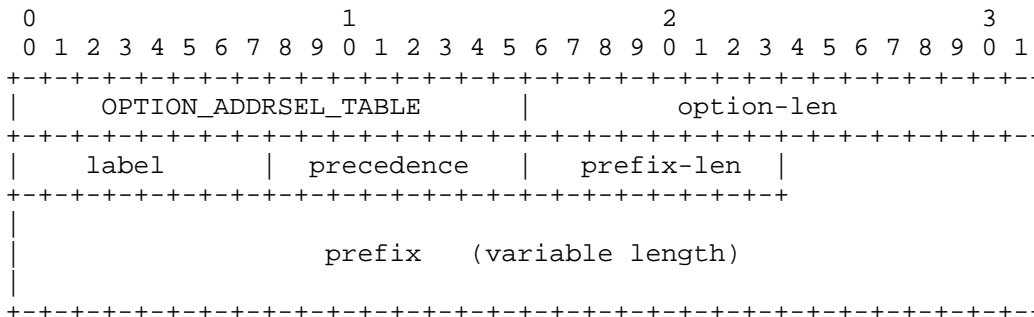


Figure 2: Address Selection Policy Table Option Format

option-code: OPTION_ADDRSEL_TABLE (85).

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. This field is used to deliver a label value in the [RFC6724] policy table.

precedence: An 8-bit unsigned integer; this value is used for sorting destination addresses. This field is used to deliver a precedence value in the [RFC6724] policy table.

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. If an option with a prefix length greater than 128 is included, the whole Address Selection option MUST be ignored.

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.

This field is padded with zeros up to the nearest octet boundary when prefix-len is not divisible by 8. This can be expressed using the following equation: $(\text{prefix-len} + 7)/8$

So, the length of this field should be between 0 and 16 bytes.

For example, the prefix 2001:db8::/60 would be encoded with a prefix-len of 60; the prefix would be 8 octets and would contain octets 20 01 0d b8 00 00 00 00.

3. Processing the Address Selection Option

This section describes how to process a received Address Selection 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. Ultimately, while the node's administrator can control how to deal with the received policy information, the implementation SHOULD follow the method described below uniformly to ease troubleshooting and to reduce operational costs.

3.1. Handling Local Configurations

[RFC6724] defines two flags (A and P) and the default policy table. Also, users are usually able to configure the flags and the policy table to satisfy their own requirements.

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

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

Choice (a) SHOULD be the default, i.e., that the policy table is not explicitly configured by the user.

3.2. Handling Stale Distributed Flags and Policy Table

When the information from the DHCP server goes stale, the flags and the policy table received from the DHCP server SHOULD be deprecated. The local configuration SHOULD be restored when the DHCP-supplied configuration has been deprecated. In order to implement this, a host can retain the local configuration even after the flags and the policy table is updated by the distributed flags and policy table.

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

3.3. Handling Multiple Interfaces

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.

The policy table is defined as a whole, so the slightest addition/deletion from the policy table brings a change in the semantics of the policy.

It also should be noted that the absence of a DHCP-distributed policy from a certain network interface should not infer that the network administrator does not care about address selection policy at all, because it may mean there is a preference to use the default address selection policy. So, it should be safe to assume that the default address selection policy should be used where no overriding policy is provided.

Under the above assumptions, we can specify how to handle received policy as follows.

In the absence of distributed policy for a certain network interface, the default address selection policy SHOULD be used. A node should use Address Selection options by default in any of the following two cases:

- 1: A single-homed host SHOULD use default address selection options, where the host belongs exclusively to one administrative network domain, usually through one active network interface.

- 2: Hosts that use advanced heuristics to deal with multiple received policies that are defined outside the scope of this document SHOULD use Address Selection options.

Implementations MAY provide configuration options to enable this protocol on a per-interface basis.

Implementations MAY store distributed address selection policies per interface. They can be used effectively on implementations that adopt per-application interface selection.

4. 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 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 [RFC3315]. 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; 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.

5. 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, Secure Neighbor Discovery (SEND), and DHCP authentication, as described in Section 21 of [RFC3315]. A commonly used alternative mitigation is to employ DHCP snooping at Layer 2.

Another threat surrounds the potential privacy concern as described in the security considerations section of [RFC6724], whereby an attacker can send packets with different source addresses to a destination to solicit different source addresses in the responses from that destination. This issue will not be modified by the introduction of this option, regardless of whether or not the host is multihomed.

6. IANA Considerations

IANA has assigned option codes to OPTION_ADDRSEL (84) and OPTION_ADDRSEL_TABLE (85) from the "DHCP Option Codes" registry (<http://www.iana.org/assignments/dhcpv6-parameters/>).

7. References

7.1. 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.
- [RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012.

7.2. Informative References

- [ADDR-SEL] Chown, T., Ed., and A. Matsumoto, Ed., "Considerations for IPv6 Address Selection Policy Changes", Work in Progress, April 2013.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC3484] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, 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

The 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. Examples

[RFC5220] gives several cases where address selection problems happen. This section contains some examples for solving those cases by using the DHCP option defined in this text to update the hosts' policy table in a network, accordingly. There is also some discussion of example policy tables in Sections 10.3 to 10.7 of RFC 6724.

B.1. Ingress Filtering Problem

In the case described in Section 2.1.2 of [RFC5220], the following policy table should be distributed when the Router performs static routing and directs the default route to ISP1 as per Figure 2. By putting the same label value to all IPv6 addresses (::/0) and the local subnet (2001:db8:1000:1::/64), a host picks a source address in this subnet to send a packet via the default route.

Prefix	Precedence	Label
::1/128	50	0
::/0	40	1
2001:db8:1000:1::/64	45	1
2001:db8:8000:1::/64	45	14
::ffff:0:0/96	35	4
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

B.2. Half-Closed Network Problem

In the case described in Section 2.1.3 of [RFC5220], the following policy table should be distributed. By splitting the closed network prefix (2001:db8:8000::/36) from all IPv6 addresses (::/0) and giving different labels, the closed network prefix will only be used when packets are destined for the closed network.

Prefix	Precedence	Label
::1/128	50	0
::/0	40	1
2001:db8:8000::/36	45	14
::ffff:0:0/96	35	4
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

B.3. IPv4 or IPv6 Prioritization

In the case described in Section 2.2.1 of [RFC5220], the following policy table should be distributed to prioritize IPv6. This case is also described in [RFC6724].

Prefix	Precedence	Label
::1/128	50	0
::/0	40	1
::ffff:0:0/96	100	4
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

B.4. ULA or Global Prioritization

In the case described in Section 2.2.3 of [RFC5220], the following policy table should be distributed, or the Automatic Row Addition flag should be set to 1. By splitting the Unique Local Address (ULA) in this site (fc12:3456:789a::/48) from all IPv6 addresses (::/0) and giving it higher precedence, the ULA will be used to connect to servers in the same site.

Prefix	Precedence	Label
::1/128	50	0
fc12:3456:789a::/48	45	14
::/0	40	1
::ffff:0:0/96	35	4
2002::/16	30	2
2001::/32	5	5
fc00::/7	3	13
::/96	1	3
fec0::/10	1	11
3ffe::/16	1	12

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