

Multiple Interfaces (Mif)
Internet-Draft
Updates: 3484bis
(if approved)
Intended status: Experimental
Expires: September 14, 2011

J. Korhonen
Nokia Siemens Networks
T. Savolainen
Nokia
March 13, 2011

Controlling Traffic Offloading Using Neighbor Discovery Protocol
draft-korhonen-mif-ra-offload-01.txt

Abstract

This specification defines an extension to IPv6 Neighbor Discovery Protocol, which allows management of IPv6 traffic offloading to IPv4 and moving IPv4 traffic away from a specific interface. The specification updates the source and destination algorithms described in RFC 3484bis.

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 September 14, 2011.

Copyright Notice

Copyright (c) 2011 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.

Table of Contents

1. Introduction	3
2. Requirements and Terminology	3
3. Problem Background	3
4. Solution	4
4.1. Neighbor Discovery Offload Option	4
4.2. Lowering the Preference of IPv6 Default Addresses	5
4.3. Lowering IPv4 Default Router Preference	6
4.4. Offload Lifetime	6
5. Router Behavior	6
6. Host Behavior	7
7. Modification to Default Address Selection	7
8. Security Considerations	8
9. IANA Considerations	8
10. References	8
10.1. Normative References	8
10.2. Informative References	8
Appendix A. Address selection examples	9
A.1. Case 1: IPv6-only cellular and IPv4-only WLAN accesses	9
A.2. Case 2: WLAN access with multiple prefixes	9
A.3. Case 3: WLAN and cellular interface with cellular's IPv4 not default route	10
A.4. Case 4: Dual-stack cellular access	10
A.5. Case 5: Dual-stack cellular and single stack WLAN	10
A.6. Case 6: Coexistence with RFC4191	11
Authors' Addresses	11

1. Introduction

This specification defines an extension to Neighbor Discovery Protocol [RFC4861], which allows management of IPv6 traffic offloading to IPv4 and moving IPv4 traffic away from a specific network connection.

The described solution is intended to be used during transition towards IPv6, during which time multi-interfaced hosts are often likely to have network interfaces with IPv4-only capability. A common scenario where coexistence of IPv4 and IPv6 network interfaces is expected to occur is when a smartphone has IPv6-enabled cellular connection and IPv4-only WLAN connection active at the same time.

This specification updates the source and destination algorithm described in RFC 3484bis [I-D.ietf-6man-rfc3484-revise]

2. Requirements and Terminology

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].

3. Problem Background

Current Internet hosts generally prefer IPv6 addresses over IPv4 addresses when performing source and destination address selections, as is recommended in [I-D.ietf-6man-rfc3484-revise].

A multi-interfaced host may have IPv6 enabled on a more 'expensive' interface and a 'cheaper' interface may have support only for IPv4. In such a scenario it might be desirable for hosts to prefer IPv4 in communication instead of IPv6.

The above mentioned problem can occur, for example, when a smartphone has simultaneously IPv6-enabled cellular connection ([I-D.korhonen-v6ops-3gpp-eps]) and IPv4-only WLAN connectivity active. When connecting to dual-stack capable destinations it would oftentimes be generally more efficient to use WLAN network interface. Furthermore, a cellular network operator may want hosts to offload traffic away from cellular network whenever hosts have alternate network accesses available.

Similar issue can arise also when a host has multiple interfaces with IPv4 connectivity. The cheaper interface should oftentimes be used for the communication, but it may not be clear for a host which one

of the available interfaces it should prefer.

4. Solution

This document introduces a new Neighbor Discovery option that a network can use to communicate 'lower-than-IPv4' preference for advertised prefix(es), and hence for host's IPv6 address, and also the level of router's willingness to act as an IPv4 default router.

The new Neighbor Discovery option was chosen to support hosts without DHCPv6 [RFC3315] support and also to work on networks not utilizing DHCPv6.

The new Neighbor Discovery option shall be phased out when IPv4 usage diminishes.

4.1. Neighbor Discovery Offload Option

This specification defines a new Neighbor Discovery [RFC4861] option called Offload (Type TBD) to be used in Router Advertisements. The option is illustrated in Figure 1. Router and hosts implementing this specification MUST understand the Offload option.




Figure 1: Router Advertisement Offload Option

Type

TBD by IANA.

Length

MUST be set to 1.

L (Lower-than-IPv4 Preference)

In addition to [RFC4191] defined Prf handling, the additional 'L' flag bit indicates 'Lower-than-IPv4' preference. From the [RFC4191] point of view, the 'Lower-than-IPv4' does not have any affect. The 'L' bit affects the source and destination address

selection for IPv6 addresses configured from prefixes advertised by the Router Advertisement containing the Offload option.

[*Discussion*: The 'L' flag has partially the same effect as setting the preferred lifetime to zero in the Prefix Information option. However, we did not want to change Prefix Information option configuration in the router as this option can then be used independently to control preferences, whether they are on or off.]

D (Default IPv4 Gateway Preference)

Indicates the willingness of the Dual-Stack capable router (who originated the Router Advertisement) to serve as a default gateway for the IPv4 traffic. If 'D' is unset (0) then the router indicates no specific to be or not to be a default gateway for IPv4 traffic. If 'D' is set (1) then the router explicitly indicates it is not willing to serve as a default gateway for IPv4 traffic if there are other usable gateways present in the same or other available interfaces.

Reserved

A 14-bit unused field. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

Gateway

The address of the dual-stack router's IPv4 interface used as the next-hop from hosts point of view for sending and receiving IPv4 traffic on this link. The IPv4 address MUST belong to the same interface that originated the Router Advertisement containing this option. If the router is IPv6 only or the 'D' bit unset (0), then this field MUST be set to unspecified address (0.0.0.0).

The behavior of 'lower-than-IPv4 Preference' (see Section 4.2) and 'Default IPv4 Gateway Preference' (see Section 4.3) are discussed in more detail in the following sections. The Offload option is only used in Router Advertisement messages.

4.2. Lowering the Preference of IPv6 Default Addresses

Router originating a Router Advertisement (RA) with the 'L' bit set in the Offload option indicates that it SHOULD NOT be used for forwarding IPv6 traffic for destinations that are also reachable with IPv4 (via other interfaces) or IPv6 destinations are also reachable using other interfaces.

If a host implements this specification, the 'L' bit required

behavior can be achieved either by host implementation dependent means (especially relevant in cellular hosts that have a per application 'silo view' of the networking stack) or by modifying the [I-D.ietf-6man-rfc3484-revise] policy table and the default address selection algorithm.

In the latter case, the [I-D.ietf-6man-rfc3484-revise] is modified so that the addresses configured from the prefixes advertised by the 'lower-than-IPv4 Preference' router MUST be treated as 'deprecated' if there is no more specific route for the intended IPv6 destination (i.e. other than default router for the intended destination would be selected). The modification concerns both the source address selection algorithm and the destination address selection algorithm. The expected behavior is that other available preferred IPv6 addresses get selected over 'lower-than-IPv4' IPv6 addresses and even IPv4 destinations are preferred over IPv6 destinations when only 'lower-than-IPv4' IPv6 addresses are available as IPv6 source addresses.

4.3. Lowering IPv4 Default Router Preference

The 'D' flag bit in the Offload option indicates the willingness of the Router Advertisement originating Dual-Stack capable router to serve as a default gateway for IPv4 traffic. When 'D' is unset (0), the router does not indicate any preference of being or not being a default gateway for IPv4 traffic. If 'D' is set (1), the router indicates that it SHOULD NOT be used as a default gateway for IPv4 traffic, if other default gateways are present in the same or other available interfaces. The 'Gateway' field in the Offload option contains the IPv4 address of the Dual-Stack interface that originated the Router Advertisement. The address serves as the identification of the next-hop IPv4 routers.

4.4. Offload Lifetime

The lifetime of the [I-D.ietf-6man-rfc3484-revise] modifications and IPv4 default gateway preferences caused by the Offload option are tied to the lifetime indicated in the Router Advertisement. Also, if the router sends a new Router Advertisement without the Offload option before the router lifetime expires, it is an indication to the receiving hosts that any existing Offload option caused state/information MUST be removed.

5. Router Behavior

A router configuration SHOULD allow network administrator to add and configure this option into Router Advertisement messages. The

configuration can be selectively enabled (the Offload option is included in the Router Advertisement) or disabled (the Offload option is not included in the Router Advertisement).

6. Host Behavior

A multi-interface capable host SHOULD monitor presence of this option in received Router Advertisement messages. When the Offload option is received, the source and destination selection algorithms defined in [I-D.ietf-6man-rfc3484-revise] shall be temporarily modified as described in Section 4.2. The IPv4 source address selection and default gateway preferences shall temporarily be updated as described in 4.3.

If the host receives a Router Advertisement without the Offload option and there is an existing state created by an earlier received Offload option, then the host MUST remove all default address selection algorithm and IPv4 default gateway preferences modifications. The removals concerns the prefixes configured from router where the router advertisement was received.

7. Modification to Default Address Selection

The 'lower-than-IPv4 Preference' affects the Source Address Selection Rule 3. The notation Lower(SA) returns true if the address SA was configured from the prefixes advertised by a 'lower-than-IPv4 Preference' router. Lower(SA) returns false if the address SA was configured from prefixes advertised by other than 'lower-than-IPv4 Preference' router. The notation Default(D) returns false if the address D has more specific routes (i.e. other than the default route). Default(D) returns true if the address D points only to a default route. The modified Rule 3 would be as follows:

Rule 3: Avoid deprecated addresses.

The addresses SA and SB have the same scope. If Lower(SA) == true and Default(D) == true, then mark SA temporarily as "deprecated". If Lower(SB) == true and Default(D) == true, then mark SB temporarily as "deprecated". If one of the two source addresses is "preferred" and one of them is "deprecated" (in the [RFC4862] sense), then prefer the one that is "preferred."

Similar modification also concerns the Destination Address Selection Rule 3 when checking whether a candidate source address for a given destination is deprecated.

8. Security Considerations

The Offload option allows malicious hosts and routers to affect a victim host's next hop and default address selection if spoofing of Router Advertisements are possible on the access link. This is a well-known and understood security threat [RFC3756] and can be mitigated using, for example, Secure Neighbor Discovery [RFC3971].

9. IANA Considerations

This specification defines a new Neighbor Discovery option described in Section 4.1.

10. References

10.1. Normative References

- [I-D.ietf-6man-rfc3484-revise]
Matsumoto, A., Kato, J., and T. Fujisaki, "Update to RFC 3484 Default Address Selection for IPv6", draft-ietf-6man-rfc3484-revise-01 (work in progress), October 2010.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.

10.2. Informative References

- [I-D.korhonen-v6ops-3gpp-eps]
Korhonen, J., Soininen, J., Patil, B., Savolainen, T., Bajko, G., and K. Iisakkila, "IPv6 in 3GPP Evolved Packet System", draft-korhonen-v6ops-3gpp-eps-06 (work in progress), February 2011.
- [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.
- [RFC3756] Nikander, P., Kempf, J., and E. Nordmark, "IPv6 Neighbor

Discovery (ND) Trust Models and Threats", RFC 3756, May 2004.

[RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.

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

Appendix A. Address selection examples

Link-local addresses are omitted in all following examples. The assumption is that possible destinations have a global scope and all IPv6 enabled interfaces have at least one global scope IPv6 address. Therefore, the default address selection would always output global scope addresses over link-local addresses.

A.1. Case 1: IPv6-only cellular and IPv4-only WLAN accesses

A host has obtained global IPv6 address, 2001:db8::2, on a cellular interface and with it has received Neighbor Discovery option with 'lower-than-IPv4' preference. The host also has global IPv4 address, 192.0.2.2, on a WLAN interface.

When connecting to a dual-stack enabled destination, both 2001:db8::2 and 192.0.2.2 are considered as source addresses candidates. IPv4 address is selected, because 2001:db8::2 is considered deprecated. Hence host uses WLAN for communication.

When connecting to IPv6-only destination, 2001:db8::2 is selected and cellular network used, as there are no other IPv6 addresses available.

A.2. Case 2: WLAN access with multiple prefixes

A host has obtained two global IPv6 addresses, one of which was from a router indicating 'lower-than-IPv4' preference. For example, 2001:db8:1::2 from router with 'lower-than-IPv4' preference and 2001:db8:2::3 from router without any special preferences.

When connecting to IPv6-only destination, both addresses are considered as source address candidates. Source address selection chooses 2001:db8:2::3 as 2001:db8:1::2 is considered deprecated (Lower(2001:db8::2) == true and Default(D) == true).

A.3. Case 3: WLAN and cellular interface with cellular's IPv4 not default route

A host has obtained IPv6 address, 2001:db8::2, and IPv4 address, 192.0.2.2, from cellular network. The network has indicated 'lower-than-IPv4' preference for IPv6 and 'not your default router' for IPv4. The host also has dual-stack WLAN access with 2001:db8:1::3 and 192.0.2.30 addresses.

When connecting to IPv4-only destination, host selects 192.0.2.30 as source address because default gateway on the interface of 192.0.2.2 address is 'not default gateway'. WLAN is used for communication.

When connecting to IPv6-only destination, host selects 2001:db8:1::3 from WLAN interface as the 2001:db8::2 is considered deprecated (Lower(2001:db8::2) == true and Default(D) == true). WLAN is used for communication.

When connecting to dual-stack destination, host selects from the four candidate addresses 2001:db8:1::3, as IPv6 is preferred in general and as that address is not deprecated. WLAN is used for communication.

A.4. Case 4: Dual-stack cellular access

A host has obtained IPv6 address, 2001:db8::2, and IPv4 address, 192.0.2.2, from cellular network. The network has indicated 'lower-than-IPv4' preference.

When connecting to a dual-stack enabled destination, both addresses are considered as candidate source addresses. IPv4 address is chosen, because IPv6 address is considered deprecated.

A.5. Case 5: Dual-stack cellular and single stack WLAN

A host has obtained IPv6 address, 2001:db8::2, and IPv4 address, 192.0.2.2, from cellular network. The network has indicated 'lower-than-IPv4' preference for IPv6 and 'not your default router' for IPv4. The host also has WLAN access with 192.0.2.30 address.

When connecting to dual-stack destination, all three addresses are considered as source address candidates. The IPv4 address from WLAN, 192.0.2.30, is selected as the IPv6 address, 2001:db8::2, is considered deprecated and as the IPv4 default route points to WLAN. Hence WLAN is used for communication.

A.6. Case 6: Coexistence with RFC4191

A host has obtained IPv6 address, 2001:db8:1::2/64 from cellular network. The network has indicated 'lower-than-IPv4' preference for IPv6 and a more specific route to 2001:db8:2::/48. The host also has IPv6 WLAN access with 2001:db8:3::3/64 address.

When connecting to 2001:db8:2::1 the host selects 2001:db8:1::2 from cellular interface as a source address, because `Lower(2001:db8:1::2) == true` and `Default(2001:db8:2::1) == false` and hence the 2001:db8:1::2 is not considered as deprecated address even though 'lower-than-IPv4' preference was advertised.

When connecting to 2001:db8:4::1 the host selects 2001:db8:3::3 from WLAN interface as a source address, because `Lower(2001:db8:2::1) == true` and `Default(2001:db8:3::3) == true` and hence 2001:db8:2::1 is considered as deprecated address.

Authors' Addresses

Jouni Korhonen
Nokia Siemens Networks
Linnoitustie 6
FI-02600 Espoo
Finland

Email: jouni.nospam@gmail.com

Teemu Savolainen
Nokia
Hermiankatu 12 D
FI-33720 Tampere
Finland

Email: teemu.savolainen@nokia.com

