Multiple Interfaces (Mif)
Internet-Draft
Intended status: Experimental
Expires: May 3, 2012

J. Korhonen
Nokia Siemens Networks
T. Savolainen
Nokia
Y. Ding, Ed.
University of Helsinki

October 31, 2011

Controlling Traffic Offloading Using Neighbor Discovery Protocol
draft-korhonen-mif-ra-offload-03.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.

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 May 3, 2012.

Copyright Notice

Table of Contents

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.

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 scenario can become a problem, for example, when a smartphone has simultaneously IPv6-enabled cellular connection ([I-D.ietf-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 interface that provides better performance at a lower price 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 the level of router's willingness to act as a router for IPv4 traffic.

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 can be used together with the Route Information option defined in [RFC4191] to communicate offloading information for specific routes.

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.

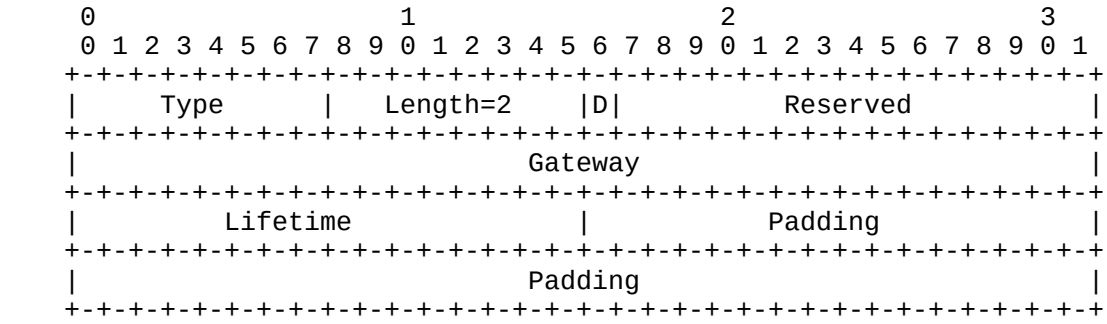

Figure 1: Router Advertisement Offload Option

Type

TBD by IANA.

Length

MUST be set to 2.

D (IPv4 Gateway Preference)

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

Reserved

A 15-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, then this field MUST be set to unspecified address (0.0.0.0) or the Neighbor Discovery Offload option MUST be omitted in all Router Advertisements originated by the router.

Lifetime

16-bit unsigned integer. The Lifetime in seconds limits the validity of state changes caused by this new option. The value of Lifetime in this option SHOULD be smaller than the value of Route Lifetime contained in the Route Information option [RFC4191], if present, in the same Router Advertisement.

Padding

The padding MUST be initialized to zero by the sender and MUST be ignored by the receiver.

The behavior of 'IPv4 Gateway Preference' (see Section 4.2) is discussed in more detail in the following sections. The usage of 'Gateway' for offloading is discussed in Section 4.3 and Section 4.4. The Offload option is only used in Router Advertisement messages.

4.2. Lowering IPv4 Router Preference

The 'D' flag bit in the Offload option indicates the willingness of Dual-Stack capable router originating the Router Advertisement to

serve as a gateway for IPv4 traffic. If 'D' is set (1), the router indicates that it SHOULD NOT be used as a gateway for IPv4 traffic, if other gateways are present in the same or other available interfaces. If 'D' is unset (0), the router does not indicate any preference of being or not being a gateway for IPv4 traffic. When 'D' is unset (0), the decision of temporarily modifying the routing status is left for hosts that receive the Offload option (see Section 4.3 and Section 4.4). 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.

4.3. IPv4 Offloading to Specific Routes

To enable offloading of IPv4 traffic to specific routes, both Offload option and Route Information option [RFC4191] MUST present in the same Router Advertisement. A host receiving such Router Advertisement need to maintain a set of status including specific route, Router Preference, and Lifetime. A specific route consists of an IPv4 gateway from the Offload option and an IPv4 prefix from the Route Information option. The Prefix field in the Route Information option SHOULD follow the IPv4-mapped IPv6 address format defined in [RFC4291]. The Prefix Length in the Route Information option is used to indicate the IPv4 prefix length. The Router Preference in the Route Information option indicates whether to prefer the IPv4 router associated with this prefix over others. The Lifetime in the Offload option determines how long the temporarily added specific route will be valid. The Lifetime field in Route Information option SHOULD be ignored.

When 'D' flag is unset (0) in the Offload option, the advertised specific route shall be added by hosts if there is no duplicated prefix matching to the advertised prefix and the advertised lifetime in Offload option is valid. If there is a matching prefix, such specific route will be updated or deleted according to the status of Lifetime and Router Preference. The Lifetime in Offload option determines whether the route will be deleted or updated depending on the existing routing status of the hosts. If the advertised Lifetime is set to 0, any matched prefix and the corresponding route MUST be removed. If Lifetime is valid, the Router Preference further determines whether the gateway of the existing route, if matched, will be substituted to the advertised one, or the lifetime for existing route will be updated.

When 'D' flag is set (1) in the Offload option, any existing specific routes with the next-hop router matching to the advertised gateway SHOULD be removed.

To avoid misconfiguration of offloading operation, only one Offload option is allowed in a single Router Advertisement.

4.4. IPv4 Offloading to Default Gateway

If there is no Route Information options containing IPv4-mapped IPv6 addresses in the same Router Advertisement, the default gateway for offloading can be added, updated, or deleted depending on the 'D' flag, Lifetime, and existing routing status on the hosts. When 'D' is set (1), the existing default gateway matching to the advertised one SHOULD be removed if there are other usable gateways present in the same or other available interfaces.

When 'D' is unset (0) and there is no default gateway present for the receiving interface, the advertised gateway with valid lifetime can be added. If the advertised gateway matches to the existing one on the host, depending on the advertised lifetime, the existing default gateway shall be updated to the advertised lifetime in Offload option or deleted if the lifetime is set 0. If there is a default gateway existing on the receiving interface, which does not match to the advertised gateway, the advertised one SHOULD be ignored.

4.5. Offload Lifetime

The lifetime in the Offload option determines the valid period of temporary routing changes including IPv4 gateway preferences and offloading of IPv4 traffic to specific routes and default gateway. 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). For specific route offloading, the prefix(es) advertised in the Route Information option SHOULD follow IPv4 mapped IPv6 address (e.g. ::ffff:1.2.3.4) as described in 4.3.

6. Host Behavior

A multi-interface capable host SHOULD monitor presence of Offload

option in received Router Advertisement messages. When the Offload option is received, the IPv4 gateway preferences and offloading to default gateway shall temporarily be updated as described in 4.2 and 4.4. Depending on the presence of Route Information in the same Router Advertisement, the offloading to specific IPv4 routes shall temporarily be updated as described in 4.3. Hosts SHOULD uses the lifetime value in the Offload option to determine the valid time of all routing changes caused by the Router Advertisement received.

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 IPv4 gateway preferences and offloading modifications from the previous Router Advertisement. The removals concerns the prefixes configured from router where the router advertisement was received.

7. 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]. The security of utilizing the Offload option is at the equal level to solution in [RFC4191].

8. IANA Considerations

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

9. References

9.1. Normative References

[I-D.ietf-6man-rfc3484-revise]
Matsumoto, A., Kato, J., Fujisaki, T., and T. Chown, "Update to RFC 3484 Default Address Selection for IPv6", draft-ietf-6man-rfc3484-revise-04 (work in progress), July 2011.

[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.

9.2. Informative References

[I-D.ietf-v6ops-3gpp-eps]
Korhonen, J., Soininen, J., Patil, B., Savolainen, T., Bajko, G., and K. Iisakkila, "IPv6 in 3GPP Evolved Packet System", draft-ietf-v6ops-3gpp-eps-08 (work in progress), September 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.

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

Appendix A. Address Selection Approach

A.1. 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 is 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.

A.2. 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.2.1. 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.2.2. 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.2.3. 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.2.4. 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.2.5. 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.2.6. 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

Yi Ding (editor)
University of Helsinki
P.O. Box 68
FI-00014 University of Helsinki
Finland

Email: yi.ding@cs.helsinki.fi