

Multiple Interfaces (Mif)
Internet-Draft
Intended status: Experimental
Expires: March 4, 2013

J. Korhonen
Nokia Siemens Networks
T. Savolainen
Nokia
A. Ding, Ed.
University of Helsinki
August 31, 2012

**Controlling Traffic Offloading Using Neighbor Discovery Protocol
draft-korhonen-mif-ra-offload-05.txt**

Abstract

This specification defines an extension to IPv6 Neighbor Discovery Protocol, which allows management of IPv4 traffic offloading for multi-interface dual-stack capable hosts 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 March 4, 2013.

Copyright Notice

Copyright (c) 2012 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 IPv4 Router Preference	7
4.3.	IPv4 Offloading to Default Gateway	7
4.4.	IPv4 Offloading to Specific Routes	8
4.5.	Offload Lifetime	8
5.	Router Behavior	8
6.	Host Behavior	9
7.	Security Considerations	9
8.	IANA Considerations	9
9.	Acknowledgements	9
10.	References	10
10.1.	Normative References	10
10.2.	Informative References	10
Appendix A.	Address Selection Approach	11
A.1.	Modification to Default Address Selection	11
A.2.	Address selection examples	11
A.2.1.	Case 1: IPv6-only cellular and IPv4-only WLAN accesses	11
A.2.2.	Case 2: WLAN access with multiple prefixes	12
A.2.3.	Case 3: WLAN and cellular interface with cellular's IPv4 not default route	12
A.2.4.	Case 4: Dual-stack cellular access	12
A.2.5.	Case 5: Dual-stack cellular and single stack WLAN	13
A.2.6.	Case 6: Coexistence with RFC4191	13
Authors'	Addresses	13

1. Introduction

This specification defines an extension to Neighbor Discovery Protocol [[RFC4861](#)], which allows management of IPv4 traffic offloading for multi-interface dual-stack capable hosts and moving IPv4 traffic away from a specific interface.

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-rfc3484bis](#)].

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 ([[RFC6459](#)]) 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 the cellular network whenever hosts have alternate network accesses available.

Similar issues 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 router's willingness to act as a default gateway for IPv4 traffic and IPv4 offloading information for specific routes.

The new Neighbor Discovery option was chosen to support hosts without DHCPv6 [RFC3315] and DHCPv4 Classless Static Route Option [RFC3442] support and also to work on networks not utilizing DHCPv6 and DHCPv4.

The Neighbor Discovery option proposed in this document SHOULD 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 and Figure 2. Routers and hosts implementing this specification MUST understand the Offload option.

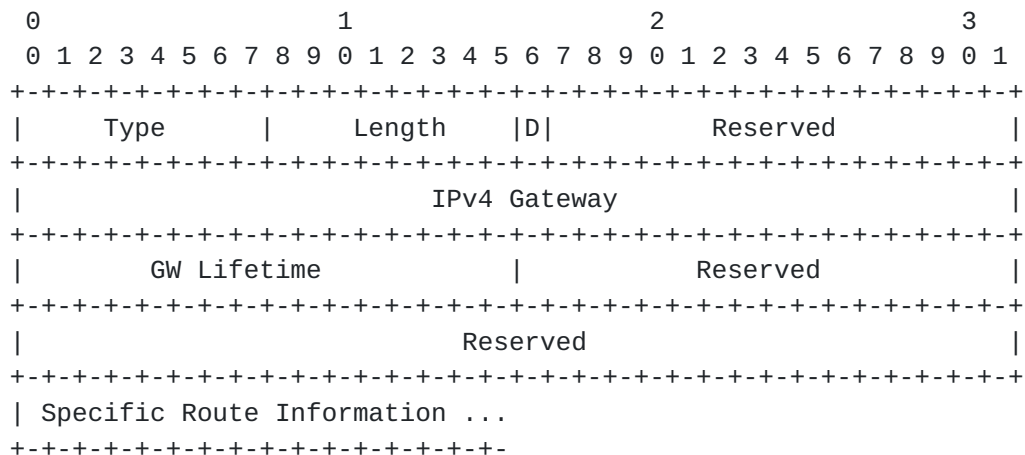


Figure 1: Router Advertisement Offload Option

Type

TBD by IANA.

Length

8-bit unsigned integer. The length of the option (including the Type and Length fields) is in units of 8 octets. The Length field depends on the optional Specific Route Information. If there is no Specific Route Information, the Length MUST be 2.

D (IPv4 Gateway Preference)

Indicates the willingness of the Dual-Stack capable router (which originated the Router Advertisement) to serve as a gateway for the IPv4 traffic. If 'D' is unset (0) then the router indicates no preference as to whether it is willing to serve as 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

Unused field. It MUST be initialized to zero by the sender and MUST be ignored by the receiver.

IPv4 Gateway

The address of a dual-stack router's IPv4 interface which is used as the next-hop from the host's 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 Offload option.

GW 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 or equal to the value of Router Lifetime contained in the header of the same Router Advertisement[RFC4191].

Specific Route Information

Optional information for IPv4 offloading to specific routes. The format is illustrated in Figure 2. An Offload option can contain multiple specific route entries.

Each Specific Route Information entry is in the length of 8 octets, containing Route Lifetime, Route Preference, Prefix Length, and IPv4 Prefix. Multiple Specific Route entries can be contained within the same Offload option.




Figure 2: Specific Route Information

Route Lifetime

16-bit unsigned integer. The length of time in seconds (relative to the time the packet is sent) that the IPv4 prefix is valid for route determination. The value of Route Lifetime SHOULD be smaller than or equal to the value of GW Lifetime contained in the same Offload option.

Prf (Route Preference)

2-bit signed integer. The preference values are encoded as follows: 01 for High; 00 for Medium (default); 11 for Low; and 10 for Reserved which MUST NOT be sent. The Route Preference indicates whether to prefer the router associated with this prefix over others, when multiple identical prefixes for different routers have been received.

Resvd (Reserved)

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

Prefix Length

8-bit unsigned integer. The number of leading bits in the IPv4 Prefix field that are valid. The value ranges from 0 to 32.

IPv4 Prefix

The IPv4 Prefix field contains an IPv4 address. The Prefix Length field contains the number of valid leading bits in the prefix. The bits in the prefix after the prefix length (if any) are

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

If the router is IPv6 only, the Neighbor Discovery Offload option MUST be omitted in all Router Advertisements originated by the router.

To avoid misconfiguration of offloading operation, a single Router Advertisement MUST contain only one Offload option.

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

[4.2.](#) Lowering IPv4 Router Preference

The 'D' flag bit in the Offload option indicates the willingness of a 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 'IPv4 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 router.

[4.3.](#) IPv4 Offloading to Default Gateway

If there is no Specific Route Information in the Offload option, the default gateway for IPv4 offloading can be added, updated, or deleted depending on the 'D' flag, GW 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 IPv4 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 lifetime of default gateway shall be updated to the advertised GW Lifetime in Offload option or deleted if the GW Lifetime is set to 0. If there is a default gateway existing on the receiving interface, which does not match the advertised gateway, the advertised one

SHOULD be ignored.

4.4. IPv4 Offloading to Specific Routes

To enable IPv4 traffic offloading to specific routes, Specific Route Information MUST be included in the same Offload option. A host receiving such Router Advertisement needs to maintain status including the IPv4 Gateway, IPv4 Prefix, Prefix Length, Route Preference, and Route Lifetime. The Route Preference in the Offload option indicates whether to prefer the IPv4 router associated with this prefix over others. The Route Lifetime in the Offload option determines how long the temporarily added specific route will be valid.

When 'D' flag is unset (0) in the Offload option, the advertised Specific Route Information shall be added by hosts if there is no duplicated IPv4 prefix matching to the advertised IPv4 prefix and the advertised Route 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 Route Lifetime and Route Preference. The Route 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 Route Lifetime is set to 0, any matched IPv4 prefix and corresponding gateway MUST be removed. If Lifetime is valid, the Route Preference further determines whether the IPv4 Gateway for the existing prefix, 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 IPv4 Gateway MUST be removed.

4.5. Offload Lifetime

The GW Lifetime in the Offload option determines the valid period of temporary routing changes including IPv4 Gateway and offloading IPv4 traffic to specific routes. If a host receives a new Router Advertisement without the Offload option, it MUST remove all existing offload state information related to the router sending the Router Advertisement.

5. Router Behavior

A router configuration SHOULD allow the 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 the presence of the Offload option in Router Advertisements received. When an Offload option is received, the IPv4 Gateway Preference and offloading status for this router shall be temporarily updated as described in 4.2 and 4.3. Depending on the presence of Specific Route Information in the same Offload option, the status of offloading IPv4 traffic to specific routes shall be temporarily updated as described in 4.4. Hosts SHOULD refer to both GW Lifetime and Route Lifetime (if present) in the Offload option to determine the valid time of 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 concern only prefixes configured from the 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. Acknowledgements

Authors would like to thank Konstantinos Pentikousis for valuable suggestions.

10. References

10.1. Normative References

- [I-D.ietf-6man-rfc3484bis]
Thaler, D., Draves, R., Matsumoto, A., and T. Chown,
"Default Address Selection for Internet Protocol version 6
(IPv6)", [draft-ietf-6man-rfc3484bis-06](#) (work in progress),
June 2012.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3442] Lemon, T., Cheshire, S., and B. Volz, "The Classless
Static Route Option for Dynamic Host Configuration
Protocol (DHCP) version 4", [RFC 3442](#), December 2002.
- [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

- [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.
- [RFC6459] Korhonen, J., Soininen, J., Patil, B., Savolainen, T.,
Bajko, G., and K. Iisakkila, "IPv6 in 3rd Generation
Partnership Project (3GPP) Evolved Packet System (EPS)",
[RFC 6459](#), January 2012.

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 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.

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

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

Email: yding@cs.helsinki.fi

