Network Working Group
Internet-Draft
Updates: 4191 (if approved)
Intended status: Standards Track
Expires: December 24, 2015

P. Pfister
Cisco Systems

June 22, 2015

Source Address Dependent Route Information Option for Router Advertisements
draft-pfister-6man-sadr-ra-01

Abstract

This document defines the Source Address Dependent Route Information option for Router Advertisements, enabling source address dependent routes to be installed in hosts by neighboring routers. It also adds a new flag to the existing Route Information option for backward compatibility purposes.

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 December 24, 2015.

Table of Contents

1. Introduction

Hosts may have multiple non-link-local addresses, possibly provided by different routers located on one or multiple links. In such situations, hosts must make sure packets with a given source address are sent to the right next-hop router. Failing in selecting the right next-hop router may, at best, induce sub-optimal routing and, at worst, cause the packet to be dropped ([RFC2827]). Rules 5 and 5.5 from the default address selection algorithm [RFC6724] make sure that, once the next-hop is chosen, care is taken to pick the right source address. Nevertheless, these rules may fail in some situations, e.g., when the same prefix is advertised on the same link by different routers. Additionally, they don't handle situations where the application picks the source-address before sending the packet.

This document defines the Source Address Dependent Route Information Option for Router Advertisements [RFC4861], enabling source address dependent routes to be installed in hosts by neighboring routers. It also adds a new flag to the Route Information Option meaning that the option may be ignored by hosts implementing this specification.

2. Source Address Dependent Route Information Option

This section defines a new Router Advertisement option called the Source Address Dependent Route Information option. Its use is similar to the Route Information option defined in [RFC4191] but also includes additional source prefix fields, allowing source address dependent routes to be installed on hosts receiving the Router Advertisement.

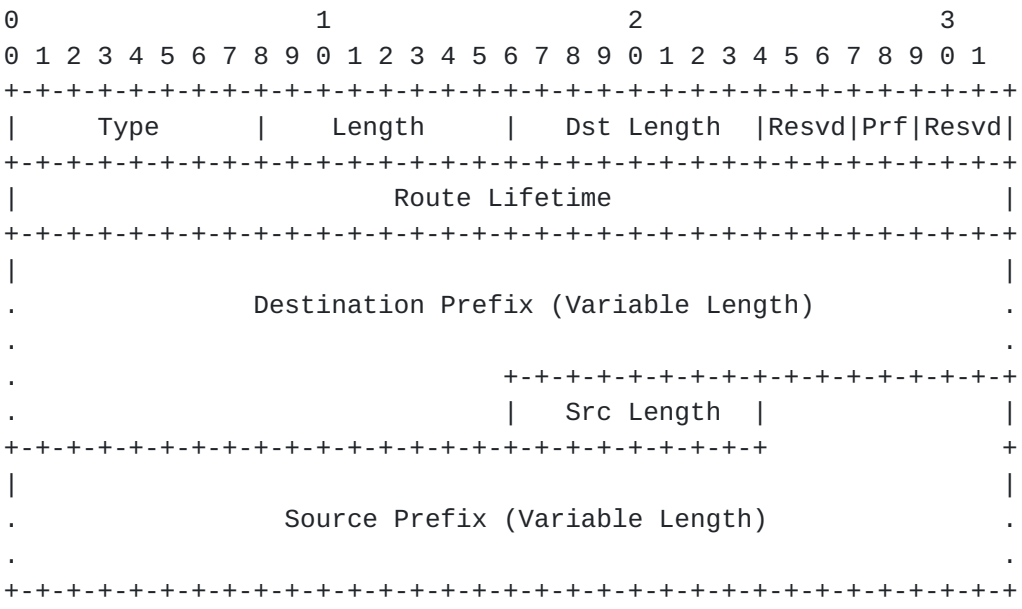

Source Address Dependent Route Information Option

Type: To be defined by IANA.

Length: The length of the option (including the Type and Length fields) in units of 8 octets. It ranges from 2 to 6.

Dst Length: The number of significant bits in the Destination Prefix field.

Resvd (Reserved): Bits reserved for futur use. They MUST be set to zero by the sender and ignored by the receiver.

Prf (Route Preference): The route preference as specified in [RFC4191]. When the Reserved value (10) is received, the option MUST be ignored.

Route Lifetime: Time in seconds (relative to the time the packet is sent) that the prefix is valid for route determination. A value of all one bits (0xffffffff) represents infinity.

Destination Prefix: The destination prefix significant bits padded to the next 8-bits boundary.

Src Length: The number of significant bits in the Source Prefix field.

Source Prefix: The source prefix significant bits padded to the next 64-bits boundary.

The following C code is given as an help for implementation:

```
#define ALIGN(bitlength, alignment) \
    (((bitlength != 0)?(((bitlength - 1) / alignment) + 1):0) * \
          (alignment / 8))

unsigned char *option;
size_t src_len_index = 8 + ALIGN(option[2], 8);
size_t total_byte_length = ALIGN((src_len_index + 1) * 8
                          + option[src_len_index], 64);
```

Note: Comments have been made regarding address alignment. There is no format providing at the same time good alignment and optimal TLV size, while aligning both source and destination prefixes would waste from 7 to 21 bytes per option. This TLV format is proposed based on implementation experience and provides both TLV size efficiency, and relative compatibility with the Route Information option (Linux implementation of this option support is less than 100 lines of code).

Comments and propositions are welcome regarding which format to adopt.

3. Route Information Option ignore flag

This document adds the Ignore flag to the Route Information option specified in [RFC4191]. It is used in order to configure type C hosts with more specific routes which will be ignored by hosts implementing this specification. Most of the time, such options with the I bit set will be used in conjunction with Source Address Dependent Route Information options including the same or a similar destination prefix.

The option is re-defined with an additional flag.

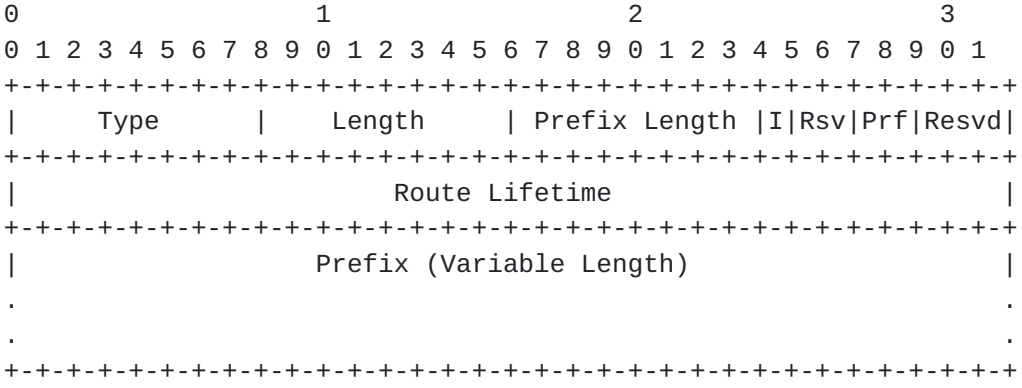

Route Information Option

I flag: Ignore flag. When this flag is set, the option MUST be ignored.

Other fields: No changes (see [RFC4191]).

4. Host Behavior

Hosts implementing this specification are referred to as type D hosts, in reference to host types A, B and C defined in [RFC4191]. As a reminder, type A hosts are hosts behaving as specified in [RFC4191]. Type B hosts behave similarly to type A hosts with the addition that they act upon the Default Router Preference values present in Router Advertisement headers. Finally, type C hosts behave as type B hosts with the addition that they act upon received Route Information Options.

This section specifies type D hosts behavior. Type D hosts MUST behave as type C hosts unless stated otherwise in this section. For the sake of clarity, in this whole section, 'host' refers to 'type D host'.

Hosts MUST use a Routing Table with source address dependent entries. Such entries have a:

- Source prefix
- Destination prefix
- Preference value
- Interface

- Next-hop router address
- Lifetime and associated timer

4.1. Selecting the next-hop router

When sending a packet, hosts MUST select the next-hop router based on the usual source address dependent routing algorithm, i.e., by picking the matching entry with, by order of precedence:

The longest destination address match.

The longest source address match.

The greatest route preference value.

In case of a tie, hosts MAY either pick one entry or use load-sharing techniques.

4.2. Receiving Source Address Dependent Route Information option

When receiving a Source Address Dependent Route Information option, a host MUST look for an existing routing entry with:

1. The same source prefix.
2. The same destination prefix.
3. The next-hop router address equal to the source address of the received Router Advertisement.
4. The outgoing interface equal to the interface the Router Advertisement is received on.

If no routing entry is found and the Route Lifetime is not null, insert a routing entry with the given source prefix, destination prefix, route preference, having as next-hop the source address of the received Router Advertisement, on the interface receiving the packet. If the Route Lifetime is not infinity, set the routing entry timer to the Route Lifetime value.

If a routing entry is found and the Route Lifetime is not null, cancel the associated timer. If the Route Lifetime is not infinity, set the timer to the Route Lifetime value. Finally, update the entry preference with the Route Preference value.

If a routing entry is found and the Route Lifetime is null, remove the routing entry.

If both destination and source prefixes specified by the option are ::/0, the router preference and route lifetime present in the option overrides the default router lifetime and default router preference present in the header of the Router Advertisement.

4.3. Receiving Route Information options

When receiving a Route Information option, a host MUST behave as follows:

If the I bit is set, ignore the option.

Otherwise, act as when receiving a Source Address Dependent Route Information option with source prefix length set to zero.

5. Router Behavior

Routers MAY send one or multiple Source Address Dependent Route Information options in their Router Advertisements.

Routers MUST NOT send multiple Route Information options with the same Prefix (no matter what the Ignore flag value is) or multiple Source Address Dependent Route Information options with the same Source and Destination Prefixes. Additionally, routers MUST NOT send a Route Information option with the Ignore bit not set and a Source Address Dependent Route Information with the source length equal to zero if the Prefix from the Route Information option is equal to the Destination Prefix from the Source Address Dependent Route Information option.

The Ignore bit is used to configure type D hosts differently from hosts of types A, B or C. Different combinations will result in different behaviors. For instance:

When injecting a source address dependent route is desired, a Source Address Dependent Route Information option is sent in every RA. Depending on the context, a Route Information with the same prefix and the Ignore bit set MAY be sent as well in order to inject a non source address dependent route into type C hosts. Obviously, Source Address Dependent Route Information options can be used to inject non-source dependent routes as well. This technique and the use of the Ignore bit allow type C hosts and type D hosts to be configured with possibly independent routes.

When injecting a non source address dependent route is desired, the router MAY either use a Route Information option with the Ignore flag not set, in which case both type C and D hosts will be configured, or use a Source Address Dependent Route Information

option with a source prefix ::/0, in which case type C hosts will not be configured.

When a Source Address Dependent Route Information option is removed from the set of advertised options, or when the interface ceases to be an advertising interface, the router SHOULD send up to MAX_INITIAL_RTR_ADVERTISEMENTS unsolicited Router Advertisements, using the same rule as in [RFC2461], with the Route Lifetime set to zero in all Source Address Dependent Route Information options that have become invalid.

6. Security Considerations

This document allows routers to configure neighboring hosts with source address dependent routing entries. Based on [RFC4191], attackers can inject default routes to type A and B hosts as well as destination address dependent routes to type C hosts. The Source Address Dependent Route Information option adds the ability for attackers to inject even more specific routes, making attacks slightly harder to detect.

7. IANA Considerations

IANA is kindly asked to reserve a Router Advertisement option type to be used by the Source Address Dependent Route Information option.

8. Acknowledgments

The author would appreciate reviews and comments.

9. References

9.1. Normative References

[RFC2461] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.

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

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007.

9.2. Informative References

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.

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

Author's Address

Pierre Pfister
Cisco Systems
Paris
France

Email: pierre.pfister@darou.fr