Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 11, 2015

B. Sarikaya
Huawei USA

December 8, 2014

IPv6 RA Options for Next Hop Routes
draft-sarikaya-6man-next-hop-ra-04

Abstract

This document proposes new Router Advertisement options for configuring next hop routes on the mobile or fixed nodes. Using these options, an operator can easily configure nodes with multiple interfaces (or otherwise multi-homed) to enable them to select the routes to a destination. Each option is defined together with definitions of host and router behaviors. This document also proposes the router advertisement extensions for source address dependent routing.

Table of Contents

1. Introduction

IPv6 Neighbor Discovery and IPv6 Stateless Address Autoconfiguration protocols can be used to configure fixed and mobile nodes with various parameters related to addressing and routing [RFC4861], [RFC4862], [RFC4191]. DNS Recursive Server Addresses and Domain Name Search Lists are additional parameters that can be configured using router advertisements [RFC6106].

Router Advertisements can also be used to configure fixed and mobile nodes in multi-homed scenarios with route information and next hop address. Different scenarios exist such as the node is simultaneously connected to multiple access network of e.g. WiFi and 3G. The node may also be connected to more than one gateway. Such connectivity may be realized by means of dedicated physical or logical links that may also be shared with other users nodes such as in residential access networks.

Host configuration can be done using DHCPv6 or using router advertisements. A comparison of DHCPv6 and RA based host

configuration approaches is presented in [I-D.yourtchenko-ra-dhcpv6-comparison].

2. 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. Default Route Configuration

A host, usually a mobile host interested in obtaining routing information usually sends a Router Solicitation (RS) message on the link. The router, when configured to do so, provides the route information using zero, one or more Next Hop Address and Route Information options in the router advertisement (RA) messages sent in response.

The route options are extensible, as well as convey detailed information for routes.

RS and RA exchange is for next hop address and route information determination and not for determining the link-layer address of the router. Subsequent Neighbor Solicitation and Neighbor Advertisement exchange can be used to determine link-layer address of the router.

It should be noted that the proposed options in this document will need a central site-wide configuration mechanism. The required values can not automatically be derived from routing tables.

Next hop address and related route information may be provided by some other means such as directly by the next hop routers. In this document we assume that next hop routers are not able to provide this information. One solution would be to develop an inter-router protocol to instigate the next hop routers to provide this information. However, such a solution has been singled out due to the complexities involved.

A non-trustworthy network may be available at the same time as a trustworthy network, with the risk of bad consequences if the host gets confused between the two. These are basically the two models for hosts with multiple interfaces, both of which are valid, but which are incompatible with each other. In the first model, an interface is connected to something like a corporate network, over a Virtual Private Network (VPN). This connection is trusted because it has been authenticated. Routes obtained over such a connection can probably be trusted, and indeed it may be important to use those routes. This is because in the VPN case, you may also be connected

to a network that's offered you a default route, and you could be attacked over that connection if you attempt to connect to resources on the enterprise network over it.

On the other, non-trustworthy network scenario, none of the networks to which the host is connected are meaningfully more or less trustworthy. In this scenario, the untrustworthy network may hand out routes to other hosts, e.g. those in the VPN going through some malicious nodes. This will have bad consequences because the host's traffic intended for the corporate VPN may be hijacked by the intermediate nodes.

Router advertisement extensions described in this document can be used to install the routes. However, the use of such a technique makes sense only in the former case above, i.e. trusted network. So the host MUST have an authenticated connection to the network it connects so that the router advertisements can be trusted before establishing routes.

4. Source Address Dependent Routing

In multihomed networks there is a need to do source address based routing if some providers are performing the ingress filtering defined in BCP38 [RFC2827]. This requires the routers to consider the source addresses as well as the destination addresses in determining the next hop to send the packet to.

The routers may be informed about the source addresses to use in routing using extensions to the routing protocols like IS-IS defined in [ISO.10589.1992] [I-D.baker-ipv6-isis-dst-src-routing] and OSPF defined in [RFC5340] [I-D.baker-ipv6-ospf-dst-src-routing]. In this document we define the router advertisement extensions for source address dependent routing.

Routing protocol extensions for source address dependent routing does not avoid a host using a source address that may be subject to ingress filtering when sending a packet to one of the next hops. In that case the host receives an ICMP source address failed ingress/ egress policy error message in which case the host must resend the packet trying a different source address. The extensions defined in this document aims at avoiding this inefficiency in packet forwarding at the host.

More information on the scenarios, their analysis and why host based approach to source address dependent routing is needed, are presented in [I-D.sarikaya-6man-sadr-overview].

5. Host Configuration

Router advertisement options defined in this document are used by Type C hosts.

As defined in [RFC4191] Type C host uses a Routing Table instead of a Default Router List.

The hosts set up their routing tables based on the router advertisement extensions defined in this document. The routes established are used in forwarding the packets to a next hop based on the destination prefix/address using the longest match algorithm. The hosts MUST keep Route Prefix that it received together with Next Hop Address, Source Address options in a stable storage. This will enable the host to consistently use these options as described next.

In case the host receives Next Hop Address with Source Address and Route Prefix option, the host uses source and destination prefix/ address using the longest match algorithm in order to select the next hop to forward the packet to.

6. Router Configuration

The router MAY send one or more Next Hop Address that specify the IPv6 next hop addresses. Each Next Hop Address may be associated with one or more Route Prefix options that represent the IPv6 destination prefixes reachable via the given next hop. Router includes Route Prefix option in message to indicate that given prefix is available directly on-link. When router sends Next Hop Address that is associated with Router Prefix option, the router MUST use Next Hop Address with Route Prefix option defined in Section 11. The Route Prefix MAY contain ::/0, i.e. with Prefix Length set to zero to indicate available default route.

The router MAY send one or more Next Hop Address options that specify the IPv6 next hop addresses and source address. Each Next Hop Address may be associated with zero, one or more Source Prefix that represent the source addresses that are assigned from the prefixes that belong to this next hop. The option MAY contain Route Prefix options that represent the IPv6 destination prefixes reachable via the given next hop as defined in Figure 4. Router includes Next Hop Address with Route Prefix option and Source Prefix in the message to indicate that given prefix is available directly on-link and that any source addresses derived from the source prefix will not be subject to ingress filtering on these routes supported by these next hops.

The router MAY send one or more Next Hop Address that specify the IPv6 next hop addresses and source address. Each Next Hop Address

option may be associated with zero, one or more Source Address that represent the source addresses that are assigned from the prefixes that belong to this next hop. The option MAY contain Route Prefix options that represent the IPv6 destination prefixes reachable via the given next hop defined in Figure 5. Router includes Next Hop Address with Source Address and Route Prefix option in the message to indicate that given prefix is available directly on-link and that the source address will not be subject to ingress filtering. For the Source Address, Source Prefix option is used with prefix length set to 128.

Each Next Hop Address may be associated with zero, one or more Source Prefix that represent the source addresses that are assigned from the prefixes that belong to this next hop. The option MAY contain Route Prefix options that represent the IPv6 destination prefixes reachable via the given next hop. Router includes Next Hop Address with Route Prefix option defined in Section 11 in the message to indicate that given prefix is available directly on-link. Next Hop Address with Route Prefix option MUST be followed by a Source Prefix option defined in Section 10 to indicate that any source addresses derived from the source prefix will not be subject to ingress filtering on these routes supported by these next hops.

In home networks, there is possibility of configuring each interface of the host using Router Advertisements sent from their next hop routers. This brings the need for a new option, Router Prefix with Source Address Option defined in Figure 6 to indicate that any source addresses derived from the source prefix will not be subject to ingress filtering on these routes supported by this router.

7. RA Packet Size and Router Issues

The options defined in this document are to be used on multi-homed hosts. A mobile host would typically have two interfaces, Wi-Fi and 3G but hosts with 3 or 4 interfaces may also exist. Configuring such hosts using the options defined in this document brings up the RA packet size issue, i.e. the packet size should not exceed the maximum transmission unit (MTU) of the link.

Total size of all options defined in this document is 160 octets. Considering that 1500 bytes is the minimum MTU configured by the vast majority of links in the Internet the hosts with 3-4 interfaces or links can be easily configured by a single router advertisement message carrying the options defined here.

The router before sending the RA SHOULD check if it fits in one frame, i.e. the size does not exceed the path MTU, the router should send a single RA message. If it does not then sending the options in

consecutive RA messages should be considered, avoiding any re-assembly issues.

The routes advertised have route lifetime values. The host considers the routes in its routing table stale when the lifetime expires. The router MUST refresh these routes periodically in order to avoid stale routing table entries in the hosts.

In some cases the mobile devices with multiple interfaces become routers. Such devices may configure their routing tables using routing protocols such as RIPng or OSPFv3 [RFC7157]. RA based approach described in this document can also be used to configure such hosts.

8. Route Prefix option

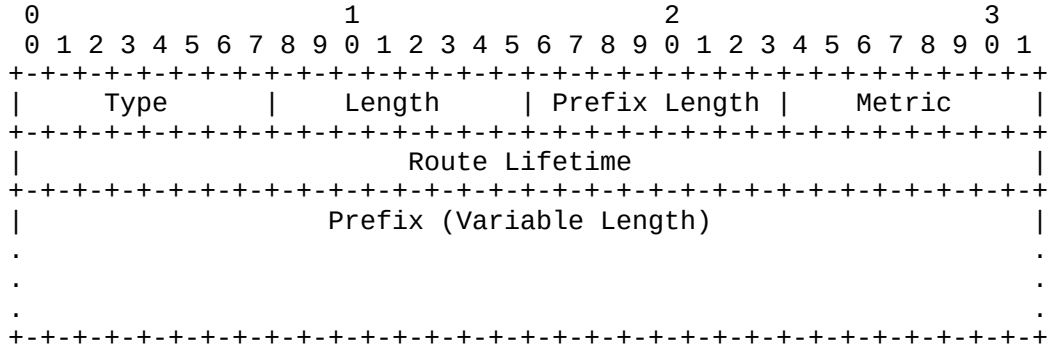

Figure 1: Route Prefix option

Fields:

Type: TBD.

Length: The length of the option (including the Type and Length fields) in units of 8 octets.

Other fields are as in [RFC4191] except:

Metric: Route Metric. 8-bit signed integer. The Route Metric indicates whether to prefer the next hop associated with this prefix over others, when multiple identical prefixes (for different next hops) have been received.

9. Next Hop Address option

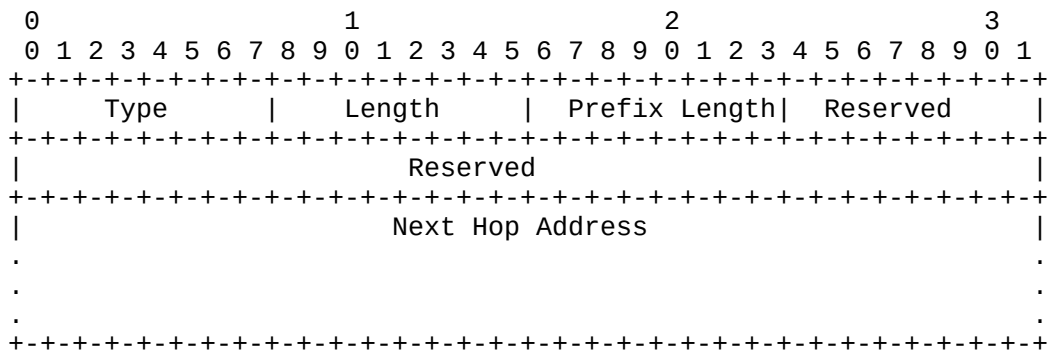

Figure 2: Next Hop Address option

Fields:

Type: TBD.

Length: The length of the option (including the type and length fields) in units of 8 octets. It's value is 3.

Prefix Length: 128

Next Hop Address: An IPv6 address that specifies IPv6 address of the next hop. It is 16 octets in length.

10. Source Address/Prefix option

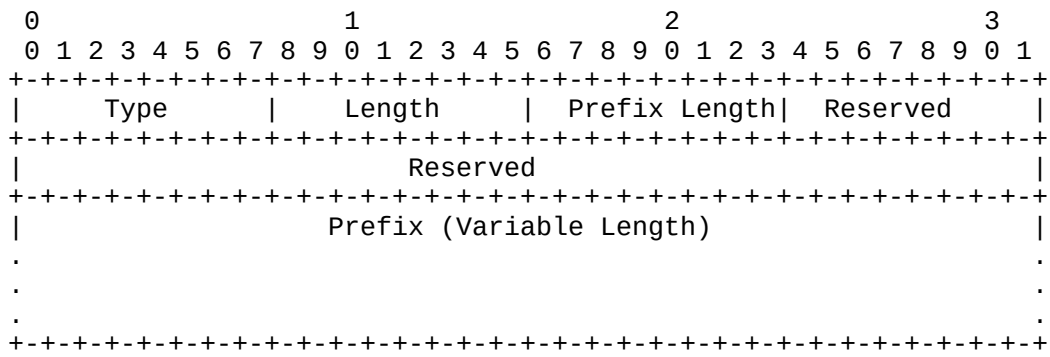

Figure 3: Source Address/Prefix option

Fields:

Type: TBD.

Length: The length of the option (including the type and length fields) in units of 8 octets. It's value is 3.

Prefix Length: An IPv6 prefix length in bits, from 0 to 128.

Prefix: An IPv6 prefix that specifies the source IPv6 prefix. It is 16 octets or less in length. Note that when the prefix length is set to 128, this option becomes a source address option.

11. Next Hop Address with Route Prefix option

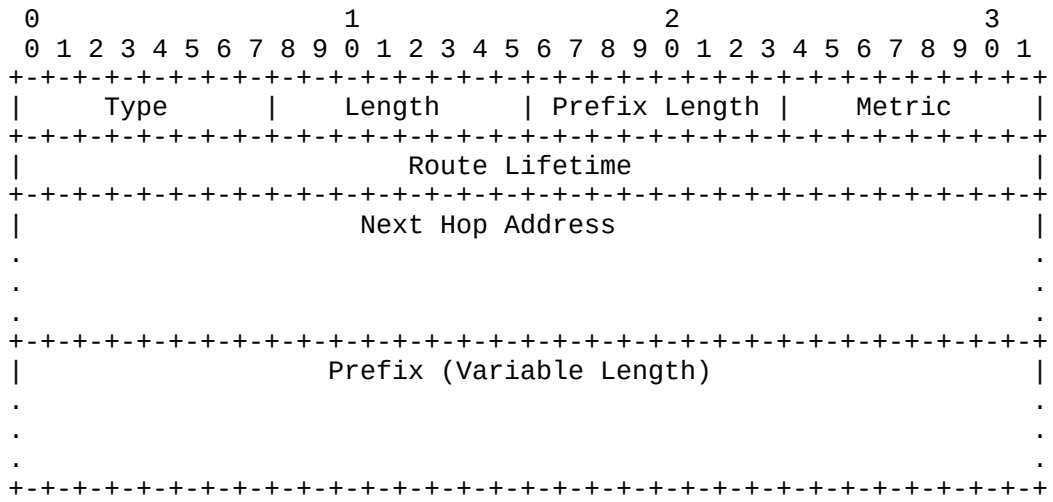

Figure 4: Next Hop Address with Route Prefix option

Fields:

Type: TBD.

Length: The length of the option (including the type and length fields) in units of 8 octets. For example, the length for a prefix of length 16 is 5.

Other fields are as in Section 8 and Section 9.

12. Next Hop Address with Source Address and Route Prefix option

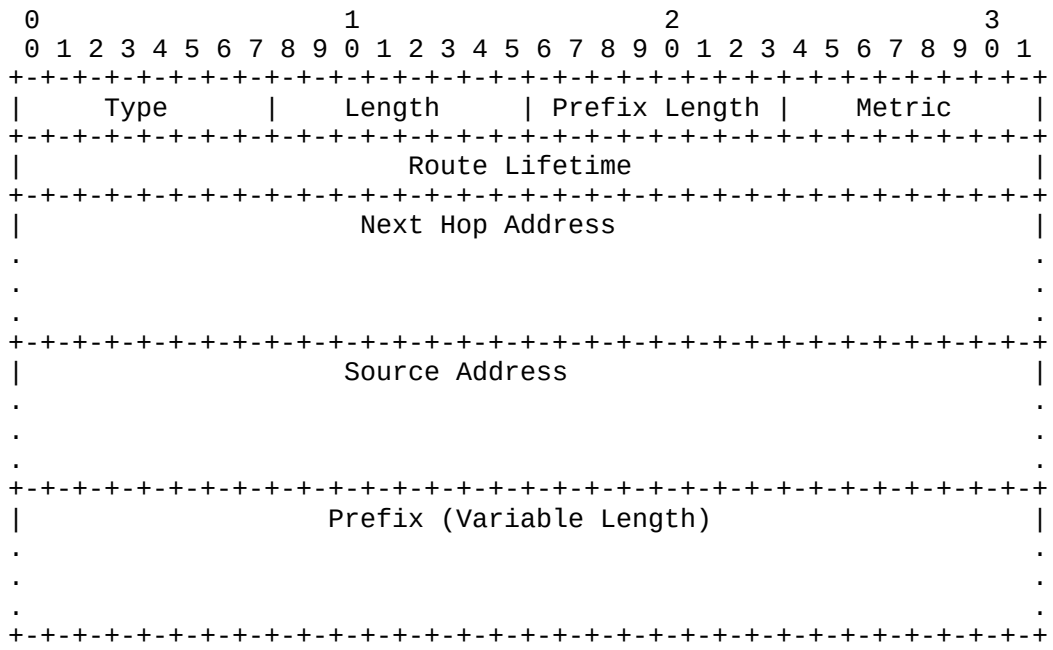

Figure 5: Next Hop Address with Source Address and Route Prefix option

Fields:

Type: TBD.

Length: The length of the option (including the type and length fields) in units of 8 octets. For example, the length for a prefix of length 16 is 7.

Other fields are as in Section 8, Section 9 and Section 10. Note that when prefix length is set to 128, the source prefix field refers to the source address.

13. Route Prefix with Source Address/Prefix Option

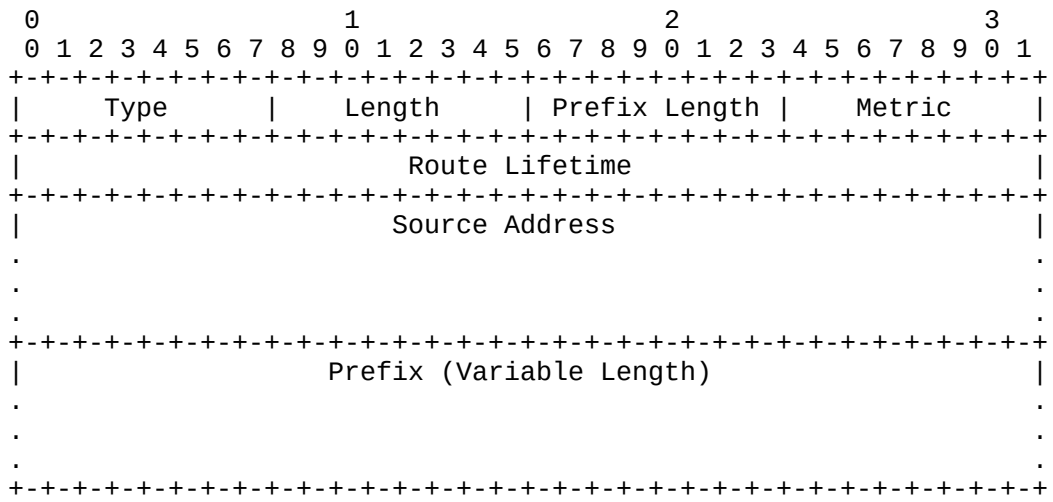

Figure 6: Route Prefix with Source Address option

Fields:

Type: TBD.

Length: The length of the option (including the type and length fields) in units of 8 octets. For example, the length for a prefix of length 16 is 5.

Other fields are as in Section 8 and Section 10.

14. Security Considerations

Neighbor Discovery is subject to attacks that cause IP packets to flow to unexpected places. Because of this, neighbor discovery messages SHOULD be secured, possibly using Secure Neighbor Discovery (SEND) protocol [RFC3971].

15. IANA Considerations

Authors of this document request IANA to assign the following new RA options:

Option Name	Type
Route Prefix	
Next Hop Address	
Source Address/Prefix	
Next Hop Address and Route Prefix	
Next Hop Address with Source Address and Route Prefix	
Route Prefix with Source Address	

Table 1:

16. Acknowledgements

Dan Luedtke, Brian Carpenter, Ray Hunter, Pierre Pfister provided many comments that have been incorporated into the document. Comments from Lorenzo Colitti, Ole Troan are much appreciated.

17. References

17.1. Normative References

[ISO.10589.1992]
International Organization for Standardization, "Intermediate system to intermediate system intra-domain-routing routine information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473), ISO Standard 10589", ISO ISO.10589.1992, 1992.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, June 1999.

[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.

[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.

[RFC4605] Fenner, B., He, H., Haberman, B., and H. Sandick, "Internet Group Management Protocol (IGMP) / Multicast Listener Discovery (MLD)-Based Multicast Forwarding ("IGMP/MLD Proxying")", RFC 4605, August 2006.

[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.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, July 2008.

[RFC7157] Troan, O., Miles, D., Matsushima, S., Okimoto, T., and D. Wing, "IPv6 Multihoming without Network Address Translation", RFC 7157, March 2014.

17.2. Informative References

[I-D.baker-ipv6-isis-dst-src-routing]
Baker, F. and D. Lamparter, "IPv6 Source/Destination Routing using IS-IS", draft-baker-ipv6-isis-dst-src-routing-02 (work in progress), October 2014.

[I-D.baker-ipv6-ospf-dst-src-routing]
Baker, F., "IPv6 Source/Destination Routing using OSPFv3", draft-baker-ipv6-ospf-dst-src-routing-03 (work in progress), August 2013.

[I-D.sarikaya-6man-sadr-overview]
Sarikaya, B., "Overview of Source Address Dependent Routing", draft-sarikaya-6man-sadr-overview-02 (work in progress), October 2014.

[I-D.yourtchenko-ra-dhcpv6-comparison]
Yourtchenko, A., "A comparison between the DHCPv6 and RA based host configuration", draft-yourtchenko-ra-dhcpv6-comparison-00 (work in progress), November 2013.

[RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, November 2010.

Author's Address

Behcet Sarikaya
Huawei USA
5340 Legacy Dr. Building 175
Plano, TX 75024

Email: sarikaya@ieee.org