Network Working Group
Internet-Draft
Updates: rfc4191, rfc4861 (if approved)
Intended status: Standards Track
Expires: July 31, 2017

F. Templin, Ed.
Boeing Research & Technology
J. Woodyatt
Google

January 27, 2017

Route Information Options in Redirect Messages
draft-templin-6man-rio-redirect-00.txt

Abstract

The IPv6 Neighbor Discovery protocol provides a Redirect function allowing routers to inform recipients of a better next hop on the link toward the destination. This document specifies a backward-compatible extension to the Redirect function to allow routers to include routing information that the recipient can associate with the next hop.

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 July 31, 2017.

Copyright Notice

Table of Contents

1. Introduction

"Neighbor Discovery for IP version 6 (IPv6)" [RFC4861][RFC2460] provides a Redirect function allowing routers to inform recipients of a better next hop on the link toward the destination. Further guidance for processing Redirect messages is given in "First-Hop Router Selection by Hosts in a Multi-Prefix Network" [RFC8028].

"Default Router Preferences and More-Specific Routes" [RFC4191] specifies a Route Information Option (RIO) that routers can include in Router Advertisement (RA) messages to inform recipients of more-specific routes. This document specifies a backward-compatible extension to allow routers to include RIOs in Redirect messages.

2. Terminology

The terminology in the normative references applies.

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]. Lower case uses of these words are not to be interpreted as carrying RFC2119 significance.

3. Route Information Options in Redirect Messages

The RIO is specified for inclusion in RA messages in Section 2.3 of [RFC4191]. The Redirect function is specified in Section 8 of [RFC4861]. This specification permits routers to include RIOs in Redirect messages so that recipients can direct future packets to a better next hop for a destination *prefix* instead of just a specific destination. This specification therefore updates [RFC4191] and [RFC4861], as discussed in the following sections.

3.1. Validation of Redirect Messages

The validation of Redirect messages follows Section 8.1 of [RFC4861], which contains the following passage:

> "The contents of any defined options that are not specified to be used with Redirect messages MUST be ignored and the packet processed as normal. The only defined options that may appear are the Target Link-Layer Address option and the Redirected Header option."

This specification updates the above statement by adding RIOs to the list of defined options that may appear.

3.2. Router Specification

The Router Specification follows Section 8.2 of [RFC4861], which provides a list of options that may appear in a Redirect message. This specification updates the list by including RIOs as permissible options. Routers therefore MAY send Redirect messages containing RIOs with values determined by a means outside the scope of this specification.

After the initial router sends Redirect messages containing RIOs that are processed by the recipient, the redirection Target MAY send its own Redirect messages containing RIOs. These Redirect messages may

be either "solicited" (i.e., an ordinary Redirect) or "unsolicited" (i.e., a Redirect generated without waiting for a packet to arrive).

An unsolicited Redirect message includes a Destination Address and Redirected Header option that are either fabricated or derived from a remembered packet that was processed at an earlier time. Alternatively, the message could omit the Redirected Header option and/or set the Destination Address field to "::" (the IPv6 unspecified address). Such a message would still satisfy the message validation checks in Section 8.1 of [RFC4861].

Any router may send RA messages with RIOs at any time, but these may be dropped along some paths over layer-2 switch fabrics that implement RA filtering.

3.3. Host Specification

The Host Specification follows Section 8.3 of [RFC4861], Section 3 of [RFC4191], and Section 3 of [RFC8028]. According to [RFC4861], a host that receives a valid Redirect message updates its destination cache per the Destination Address and its neighbor cache per the Target Address. According to [RFC4191], hosts can be classified as Type "A", "B" or "C" based on how they process valid RA messages, where a Type "C" host updates its routing table per any RIO elements included in the message. Finally, according to [RFC8028], a Type "C" host operating on a Multi-Prefix Network with multiple default routes can make source address selection decisions based on information in its routing table decorated with information derived from the source of the RIO element.

In light of these considerations, this document introduces a new Type "D" behavior for hosts with the same behavior as a Type "C" host, but which also process RIO elements in Redirect messages. Type "D" hosts process Redirect messages with RIO elements by updating 1) their neighbor cache per the Target Address, 2) their destination cache per the Destination Address, and 3) their routing tables per any RIO elements present. The host can then make source address selection decisions per [RFC8028] the same as described above.

When a Type "D" host processes a Redirect message, it SHOULD first

test the path to the Target using Neighbor Unreachability Detection (NUD) while continuing to send packets via the router that issued the Redirect until the NUD procedure converges. Thereafter, if a Route Lifetime expires (or if an RIO with Route Lifetime 0 arrives) the host removes the corresponding Prefix from its routing table and allows future packets to follow a different route.

The behaviors of Type "A", "B" and "C" hosts defined in [RFC4191] are not changed by this specification. This specification updates Section 3 of [RFC4191] by introducing a new host Type "D", and updates Section 8.3 of [RFC4861] by permitting RIOs to appear in Redirect messages.

3.3.1. Type "D" Hosts with Delegated Prefixes

Type "D" hosts may be holders of entire IPv6 prefix delegations instead of just a singleton address. For example, the host may connect an entourage of "Internet of Things" devices that derive their addresses from a delegated prefix. In that case, the host may

itself serve as a redirection target in a manner consistent with the Router Specification above. Such Type "D" hosts act like a host in terms of processing received Redirects and act like a router in terms of sending Redirects.

4. Implementation Status

The Redirect function and RIOs are widely deployed in IPv6 implementations.

5. IANA Considerations

This document introduces no IANA considerations.

6. Security Considerations

Security considerations for Redirect messages that include RIOs are the same as for any IPv6 ND messages as specified in Section 11 of [RFC4861]. Namely, the protocol must take measures to secure IPv6 ND messages on links where spoofing attacks are possible.

A spoofed Redirect message containing no RIOs could cause corruption

in the host's destination cache while a spoofed Redirect message containing RIOs could corrupt the host's routing tables. While the latter would seem to be a more onerous result, the possibility for corruption is unacceptable in either case.

"IPv6 ND Trust Models and Threats" [RFC3756] discusses spoofing attacks, and states that: "This attack is not a concern if access to the link is restricted to trusted nodes". "SEcure Neighbor Discovery (SEND)" [RFC3971] provides one possible mitigation for other cases.

[RFC6105] describes a layer-2 filtering technique called "RA Guard" intended for network operators to use in protecting hosts from receiving RA messages sent by nodes that are not among the set of default routers regarded as legitimate by the network operator. However, the RA Guard function defined in [RFC6105] does not filter ND Redirect messages. On networks with such RA Guard functions, blocked routers can use ND Redirect messages to inform hosts of routes for specific destination addresses. This draft introduces a new method by which such routers can inform Type D hosts of routes for more specific destination prefixes as well as addresses. On networks with layer-2 filters that protect hosts by restricting the delivery to hosts of both RA messages and ND Redirect messages from a limited set of legitimate routers, the host routing tables of both Type C and Type D hosts are protected by that layer-2 filtering function.

7. Acknowledgements

Joe Touch suggested a standalone draft to document this approach in discussions on the intarea list. The work was subsequently transferred to the 6man list, where the following individuals provided valuable feedback: Mikael Abrahamsson, Zied Bouziri, Brian Carpenter, Steinar Haug, Christian Huitema, Tomoyuki Sahara.

8. References

8.1. Normative References

[RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <http://www.rfc-editor.org/info/rfc791>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <http://www.rfc-editor.org/info/rfc2460>.

[RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <http://www.rfc-editor.org/info/rfc4191>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, DOI 10.17487/RFC4861, September 2007, <http://www.rfc-editor.org/info/rfc4861>.

8.2. Informative References

[RFC3756] Nikander, P., Ed., Kempf, J., and E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", RFC 3756, DOI 10.17487/RFC3756, May 2004, <http://www.rfc-editor.org/info/rfc3756>.

[RFC3971] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, DOI 10.17487/RFC3971, March 2005, <http://www.rfc-editor.org/info/rfc3971>.

[RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, DOI 10.17487/RFC6105, February 2011, <http://www.rfc-editor.org/info/rfc6105>.

[RFC8028] Baker, F. and B. Carpenter, "First-Hop Router Selection by Hosts in a Multi-Prefix Network", RFC 8028, DOI 10.17487/RFC8028, November 2016,

<http://www.rfc-editor.org/info/rfc8028>.

Appendix A. Link-layer Address Changes

Type "D" hosts send unsolicited Neighbor Advertisements (NAs) to announce link-layer address changes per standard neighbor discovery [RFC4861]. Link-layer address changes may be due to localized factors such as hot-swap of an interface card, but could also occur during movement to a new point of attachment on the same link.

Appendix B. Interfaces with Multiple Link-Layer Addresses

Type "D" host interfaces may have multiple connections to the link; each with its own link-layer address. Type "D" nodes can therefore include multiple link-layer address options in Redirects and other IPv6 ND messages. Neighbors that receive these messages can cache and select link-layer addresses in a manner outside the scope of this specification.

Authors' Addresses

Fred L. Templin (editor)
Boeing Research & Technology
P.O. Box 3707
Seattle, WA 98124
USA

Email: fltemplin@acm.org

James Woodyatt
Google
3400 Hillview Ave
Palo Alto, CA 94304
USA

Email: jhw@google.com