Internet Engineering Task Force
Internet-Draft
Updates: 4861 (if approved)
Intended status: Standards Track
Expires: 20 March 2026

M. Smith

16 September 2025

ICMPv6 Prefix Redirect Messages
draft-smith-6man-icmpv6-prefix-redirect-07

Abstract

The existing IPv6 ICMPv6 Redirect Message informs a host of a better next hop for a single destination IPv6 address. There are use cases for informing a host of a better next hop for a prefix or range of IPv6 addresses that includes or covers the single destination address that triggered the ICMPv6 redirect message. This memo specifies an ICMPv6 Prefix Redirect Message for this purpose.

Table of Contents

1. Introduction

[RFC9663] describes a method of assigning client hosts a prefix or range IPv6 addresses via DHCPv6-PD [RFC8415]. The IPv6 prefix size expected to be delegated to client hosts is a /64.

When packets are sent between hosts on the same link with different host prefixes, from and to addresses within the delegated prefixes, the sending host will normally send the packets to a default router for delivery, as the sending host is not aware that the destination address is within a prefix that is directly reachable via another host attached to the same link.

[RFC9663] advises that routers SHOULD send an ICMPv6 Redirect Message [RFC4861] to the packet sending host to inform it that the destination address of the packet is directly reachable via another host attached to the same link.

The major drawback of using existing ICMPv6 Redirect Messages in this case is that the ICMPv6 Redirect Message only redirects packets for a single destination address. Should the same sending host send a

packet to a different destination within the same destination prefix assigned to the same on-link host, it will again send that packet to a default router and the default router will again generate an ICMPv6 Redirect Message for the different destination address to the same on-link destination host.

In the scenario described by [RFC9663], a default router is aware of the prefix assigned to a host that includes the destination address that will trigger an ICMPv6 Redirect Message. Consequently, rather than generating an ICMPv6 Redirect Message for an individual destination address, it would be preferable if an ICMPv6 redirection message could convey redirection for a prefix covering a range of destination addresses assigned to a host.

Another potential use case is in a broadband access network. Similar to the [RFC9663] scenario, broadband subscribers' CPE are delegated IPv6 prefixes, such as a /48 or /56, for use on downstream LAN interfaces via DHCPv6-PD [RFC8415]. Traffic between subscribers' delegated prefixes will follow a default route to the upstream Broadband Network Gateway [BNG] router. If the subscribers' CPE are attached to the same subnet and link, the traffic between subscribers' delegated prefixes could be sent directly between their CPEs, rather than going via the upstream BNG.

A subscriber's CPE would need to be informed of the other subscriber's delegated prefix that is directly reachable over the common subnet and link. This could be achived with an ICMPv6 redirection message that conveys redirection for a prefix. Note that routers, of which subscriber CPE are an example, do not normally process ICMPv6 redirection messages [RFC4861]. [RFC7084] CPE perform some host processing of IPv6 messages, such as sending Router Soliciations and processing received Router Advertisements [RFC4861]. [RFC7084] would need to be updated to allow CPE to receive and process ICMPv6 redirection messages for a prefix.

More generally, with IPv6's vastly increased number of addresses over IPv4's, there is an increased likelihood that an IPv6 host will communicate with more IPv6 addresses. A redirect message covering an IPv6 prefix will be more efficient than the traditional redirect message for a single IPv6 address when a host is communicating with multiple IPv6 addresses within the same prefix.

This memo enhances the existing ICMPv6 Redirect Message so that it can convey an IPv6 prefix that includes the single IPv6 destination address that triggered the redirection. This enhanced redirect message is known as an ICMPv6 Prefix Redirect Message. The ICMPv6 Prefix Redirect Message is backwardly compatible with host implementations that only understand the existing single destination IPv6 address ICMPv6 Redirect Message.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. ICMPv6 Prefix Redirect Message Format

The enhanced ICMPv6 Redirect Message Format is as per [RFC4861] 4.5, with the following modification:

* The first 8 bits of the Reserved field are designated as the unsigned 8 bit integer "Prefix Length" field, carrying the length of the prefix being redirected.

* The second 8 bits of the Reserved field remain reserved, meaning set to all zeros upon transmission and ignored upon receipt.

* The final 32 bits of the Reserved field are designated as an unsigned 32 bit integer specifying the Prefix Redirect Lifetime, specifying the length of time in seconds that the prefix redirection remains valid relative to when the ICMPv6 Prefix Redirect Message is received. A value of all one bits (0xffffffff) represents infinity. A value of 0 seconds is invalid.

Note that the Destination Address field continues to carry the IPv6 address of the packet that triggered the enhanced ICMPv6 Prefix Redirect Message, retaining backward compatibility with hosts that do not understand this ICMPv6 Prefix Redirect Message.

3. Router Processing

A router implementation that follows this specification maintains two system variables per interface:

* A system varliable that specifies the largest size of a prefix that can be redirected for the interface, with a default value of 48, corresponding to redirecting an IPv6 /48 prefix or longer. The per-interface system variable’s valid range is 0 through 128.

* A system variable that specifies the the unsigned 32 bit Prefix Redirect Lifetime value in seconds used for ICMPv6 Prefix Redirect Messages sent out of this interface. The default value for this variable is 3600 seconds (or 1 hour). The valid range for this variable is 0x00000001 through 0xffffffff, with the latter specifying an infinite lifetime.

A router that can send ICMPv6 Prefix Redirect messages follows the Router Specification in section 8.2 of [RFC4861], with the additional steps:

* If the Destination Address field of the invoking IP packet is on-link, send a conventional [RFC4861] ICMPv6 Redirect Message for a single Destination Address.

* The Prefix Length field for the ICMPv6 Prefix Redirect Message is set to the length of the prefix in the route table that best matches the destination IPv6 address that triggered the prefix redirection, where the matching prefix’s length is greater than or equal to the per-interface largest size prefix system variable’s current value. If the length of the best matching route is a smaller value than the per-interface largest size prefix system variable’s value, the ICMPv6 Redirection Message Prefix Length field value set to is set to 128 (this is the ICMPv6 Prefix Redirect Message equivalent of a traditional ICMPv6 Redirect Message for a single IPv6 destinaton address. Setting the Prefix Length field value to 128 indicates that a prefix redirection could have occured if the length of the prefix in the route table was longer than the per-interface largest size prefix variable value.).

* The Prefix Redirect Lifetime field of the ICMPv6 Prefix Redirect Message is set to the interface specific lifetime value.

* The ICMPv6 Prefix Redirect Message is sent to the host that invoked the redirect, subject to ICMPv6 redirect message rate limiting [RFC4861].

A router implementation should also provide a mechanism via an administrative interface to suppress ICMPv6 Prefix Redirection Messages for a set of specified prefixes on a per-interface basis. The set of specified prefixes should support both an exact match on a prefix for which the redirection is suppressed, as well as matching

against an aggregate prefix that matches on the prefix or any matching prefixes that are longer and more specific than the specified aggregate prefix. These exact match or aggregate match prefixes should support all IPv6 prefixes and lengths ranging from /0 through /128 (for example, an aggregate prefix of ::/0 would mean suppress all ICMPv6 Prefix Redirect Messages on the interface).

4. Host Processing

A host that receives an ICMPv6 Prefix Redirect Message initially validates the message according to the steps specified in [RFC4861], Section 8.1.

4.1. Legacy Hosts

Once the message has been validated, a legacy host that does not understand the ICMPv6 Prefix Redirect message will ignore the Prefix Length and Prefix Redirect Lifetime fields because they are utilising part of the existing Reserved field, which is a backward-compatible change; [RFC4861]:

"The contents of the Reserved field, and of any unrecognized options, MUST be ignored. Future, backward-compatible changes to the protocol may specify the contents of the Reserved field or add new options; backward-incompatible changes may use different Code values."

A legacy host will process the ICMPv6 Prefix Redirect message as though it was for a single destination address, the address held in the ICMP Prefix Redirect message Destination Address field, per [RFC4861] section 8.3.

4.2. ICMPv6 Prefix Redirect Aware hosts

A host implements the following interface specific system variables:

* A system varliable that specifies the largest size of a prefix that will be accepted in an ICMPv6 Prefix Redirect Message, with a default value of 48, corresponding to redirecting an IPv6 /48 prefix or longer. The per-interface system variable’s valid range is 0 through 128.

* A system variable that specifies the maximum value of the the unsigned 32 bit Prefix Redirect Lifetime value in seconds that will be accepted in an ICMPv6 Prefix Redirect Messages. The default value for this variable is 3600 seconds (or 1 hour). The valid range for this variable is 0x00000001 through 0xffffffff, with the latter specifying an infinite lifetime.

A host that implements this specification initially validates the ICMPv6 Prefix Redirect Message according to [RFC4861] section 8.1. A host implementing this specification then performs the following steps. If any of the message validation steps fail, the ICMPv6 Prefix Redirect Message is silently discarded, or is discarded and a suitable ICMPv6 error counter incremented :

* The Prefix Length field value is no less than a per-interface system implementation variable that limits how large the redirected prefix can be.

* The Prefix Redirect Lifetime field value is greater than 0, and less than or equal to the interface specific maximum accepted lifetime value.

* A host then combines the ICMPv6 Prefix Redirect Message Destination Address field with the Prefix Length field to determine the prefix that is being redirected.

* An ICMPv6 Prefix Redirect Aware host will then update its route table with a route for the redirected prefix information and the ICMPv6 Prefix Redirect Message Target Address as the route’s next hop.

Over time, the host decrements the lifetime of the redirect prefix route, unless the received Prefix Redirect Lifetime value was infinity. This redirected prefix route MUST be removed from the host’s route table if either the next hop becomes unreachable, as detected by Neighbor Unreachability Detection (NUD) [RFC4861], regardless of the remaining redirected prefix route lifetime, or the lifetime of the redirected prefix route reaches 0 seconds.

5. Incompatible Hosts

It is possible that some hosts may be incompatible with this specification, because they validate and possibly interpret values in the existing ICMPv6 Redirect Message reserved field. There are two possible scenarios in this case.

Firstly, a host may, rather than ignoring the reserved field, per [RFC4861] instructions, expect it to be set to all zeros, perhaps for security reasons. In this case it may ignore the ICMPv6 Prefix Redirect Message with non-zero values in the formerly reserved field. While this wouldn’t prevent routing from working, packets from this host would continue to follow a less than optimal forwarding path, and would continue to trigger ICMPv6 Prefix Redirect Messages due to the less than optimal forwarding path.

In the second scenario, a host and cooperating router implementation may have used the reserved field for some other non-public, out-of-specification function. The values stored in the previously reserved fields in the ICMPv6 Prefix Redirection Message may interfere with this function. Although possible, this seems fairly unlikely, since the existing ICMPv6 Redirect Message has a single and definite purpose. An out of specification use of the reserved field would be quite an abuse of the purpose of the existing ICMPv6 Redirect Message.

Should either of these situations occur, their effects can be mitigated or avoided by (a) disabling ICMPv6 Prefix Redirection Messages on the link to which these imcompatible hosts are attached, at the cost of more optimal forwarding for hosts that are compatible with ICMPv6 Prefix Redirect Messages, or (b) moving the incompatible hosts to a different link where ICMPv6 Redirect Messages aren’t enabled, leaving the remaining hosts on the original link to benefit from issued ICMPv6 Prefix Redirect Messages.

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

The security threats for ICMPv6 Prefix Redirect Messages are the same as those for the ICMPv6 Redirect Message; see [RFC4861].

The significant security difference between an ICMPv6 Prefix Redirect Message and an ICMPv6 Redirect Message is that a successful malicious ICMPv6 Prefix Redirect Message will redirect traffic for a range or a prefix’s worth of IPv6 addresses rather than just a single address. By default, a prefix redirection is limited to a /48. The size of the prefix that is accepted in an ICMPv6 Prefix Redirect Message can be reduced to a smaller prefix, such as a /56 or a /64, if necessary.

To further mitigate the impact of a malicious prefix redirect, the ICMPv6 Prefix Redirect Message contains a Prefix Redirect Lifetime field, which defaults to 3600 seconds or 1 hour, which limits how long a malicous ICMPv6 Prefix Redirect takes effect. The Prefix Redirect Lifetime field value accepted by a host can be reduced if necessary.

8. References

8.1. Normative References

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

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

8.2. Informative References

[exampleRefMin] Surname [REPLACE], Initials [REPLACE]., "Title [REPLACE]", 2006.

[exampleRefOrg] Organization [REPLACE], "Title [REPLACE]", 1984, <http://www.example.com/>.

Acknowledgements

Thanks to Brian Carpenter for his review and comments, including pointing out the issue of incompatible host implementations that may validate or interpret the reserved field.

Author’s Address

Mark Smith
PO BOX 521
Heidelberg Victoria 3084
Australia
Email: markzzzsmith@gmail.com