Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 15 June 2026

M. Smith

12 December 2025

Nearly IPv6 Only Dual Stack Hosts
draft-smith-v6ops-nearly-ipv6-only-dualstack-hosts-02

Abstract

Once all of a dual stack IPv4 and IPv6 host's applications support IPv6, IPv4 connectivity becomes unnecessary. However, ceasing to provide a dual stack host with an IPv4 address via a DHCPv4 server can cause the host to continue to make periodic and unsatisfied DHCPv4 requests, imposing unnecessary broadcast traffic on an attached link. In this situation, it would be useful to continue to provide a host with an IPv4 address via DHCPv4, yet prevent it from using the IPv4 address for any connectivity outside of the host itself. This memo describes a DHCPv4 method to achieve this.

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 https://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 15 June 2026.

Copyright Notice

Table of Contents

1. Introduction

Once all of a dual stack IPv4 and IPv6 host’s applications support IPv6, IPv4 connectivity becomes unnecessary. However, ceasing to provide a dual stack host with an IPv4 address via a DHCPv4 [RFC2131] server can cause the host to continue to make periodic and unsatisfied DHCPv4 requests, imposing unnecessary broadcast traffic on an attached link.

One method to satisfy the IPv4 addressing requirement is for a host use the DHCPv4 IPv6-Only Preferred option, option value 108 [RFC8925]. Providing this option in the DHCPv4 Parameter Request List indicates to the DHCPv4 server that the host doesn’t require an IPv4 address. If supported and enabled, the DHCPv4 server can inform the host via a response option 108 how long the host should suppress IPv4.

The benefit of the DHCPv4 IPv6-Only Preferred option is that it doesn’t consume IPv4 addresses from the DHCPv4 server’s address pool. However, the drawback of the DHCPv4 IPv6-Only Preferred option is that it requires hosts and their DHCPv4 servers to support this relatively new DHCPv4 option.

It would be useful to be able to prevent a non-DHCPv4 IPv6-Only Preferred option supporting dual stack legacy host from using IPv4 to reach external destinations (both on-link and off-link), while still satisifying its DHCPv4 requests for an IPv4 address. This memo describes a method of how to achieve this using traditional DHCPv4 mechanisms.

Note that this method could be used on a per-host basis, via DHCPv4 client specific options, or for all hosts attached to a link. It can be used when the DHCPv4 server is on the same link as the DHCPv4 client hosts, or via a DHCPv4 relay.

This method can be deployed in conjunction with the DHCPv4 IPv6-Only Preferred option, so that only legacy hosts will consume IPv4 addresses from the DHCPv4 pool.

2. Supressing On-Link IPv4 Connectivity

Unlike IPv6 [RFC5942], IPv4 uses an address’s subnet mask or prefix length to determine if a destination IPv4 address should be directly reachable on a link, or if a router should be used to attempt to reach the destination IPv4 address [RFC1122].

For example, a subnet mask of 255.255.255.0 (/24) for the address 203.0.113.32 assigned to an interface tells a host that the range of addresses from 203.0.113.1 through 203.0.113.31, and from 203.0.113.33 to 203.0.113.254 (.0 and .255 not being valid host addresses) should be reachable directly via the link attached to the interface. The host would then trigger an ARP Request/ARP Reply [RFC826] transaction to resolve the destination IPv4 address into a link-layer address, should address resolution be necessary.

If a host was instead provided with a 255.255.255.255 subnet mask for an address configured on a interface, the host would conclude that there are no directly reachable remote addresses on the attached link. The host would attempt to use a router to reach any and all remote IPv4 destinations because it considers all destinations to be off-link, should the host have any routers available.

If the host does not have any routers available, then the only IPv4 addresses that the host can reach will be addresses within 127.0.0.0/8 [RFC1122], and only IPv4 addresses assigned to its interfaces with a 255.255.255.255 subnet mask. Connectivity to any other IPv4 address will immediately fail, and an ICMPv4 [RFC792] Destinaton Unreachable, Host Unreachble Message will be generated by the host and passed up to the application. This should cause the application to attempt to use IPv6 to reach an equivalent destination, should IPv6 not have been previously attempted before IPv4 [RFC6724].

When using a DHCPv4 server to provide the host with an IPv4 address, the Subnet Mask option, option value 1, is used to provide a 255.255.255.255 (/32) subnet mask.

Note that as the hosts on the link with 255.255.255.255 subnet masks will not attempt to communicate with any other IPv4 addresses on the link, it would appear that a single IPv4 address could be shared across all of the 255.255.255.255 subnet mask hosts. However, the DHCPv4 server will still need to provide each host with a unique IPv4 address within the link, meaning that a pool of IPv4 addresses is still needed on the DHCPv4 server. This is because hosts may perform IPv4 Address Conflict Detection (ACD) [RFC5227] for addresses acquired via DHCPv4. Should an address conflict be detected, the host would decline to accept the IPv4 address via a DHCPDECLINE message to the DHCPv4 server [RFC2131]. It will then request another IPv4 address from the DHCPv4 server, receive the same conflicting duplicate address, and end up in a loop of continuously being provided with an IPv4 address that will continuously fail ACD.

2.1. Suppressed On-Link Connectivity Address Pool

As the hosts with 255.255.255.255 subnet masks do not attempt to reach any off-link IPv4 destinations, the IPv4 addresses in the pool on the DHCPv4 server do not need to be routable over the network or the Internet. If the DHCPv4 server is attached to the same link as the DHCPv4 clients, it would be possible to use the same address pool on different links within the network.

The IPv4 link-local prefix of 169.254.0.0/16 [RFC3927] could be a good candidate prefix to use for this 255.255.255.255 subnet mask address pool, although it would imply that link-local IPv4 connectivity is available and can be used (the 255.255.255.255 subnet mask or /32 prefix length would however show otherwise).

A pool of addresses from within the [RFC1918] IPv4 private addresses allocations are an alternative. It might be useful to designate a specific subnet or subnets for this purpose that are significantly

different from any other [RFC1918] addressing being used in the network, to aid with identifying these 255.255.255.255 hosts if necessary.

If the 255.255.255.255 subnet mask clients reach the DHCPv4 server via a link attached DHCPv4 Relay, then the DHCPv4 pool of addresses will need to be unique across the network, as the DHCPv4 Relay's IPv4 address is used via the 'giaddr' field [RFC2131] by the DHCPv4 server to select the pool of addresses to use for the relayed DHCPv4 requests.

3. Supressing Off-Link IPv4 Connectivity

This is simply and easily achieved by not providing the host with any routers, by not providing the DHCPv4 client the Router Option, option value 3.

Note that any routers attached to the link do not need an IPv4 address if all hosts are being provided with the 255.255.255.255 subnet mask, unless the router is acting as a DHCPv4 Relay, as discussed previously.

4. Other DHCPv4 Options

Since a host with a 255.255.255.255 subnet mask's IPv4 connectivity is very limited, there is no point or purpose in providing any DHCPv4 options to the host other than the Subnet Mask Option and the Router Option.

Other host configuration parameters, such as DNS resolver IPv6 addresses, are supplied via IPv6 host configuration mechanisms [RFC6106][RFC8415].

5. Lease Times

Since the hosts with 255.255.255.255 subnet masks aren't using IPv4 to reach external destinations, and the DHCPv4 pool can be large since it is from a local or private IPv4 address space, IPv4 address lease times can be significantly bigger than typical e.g. 24 hours or even a number of weeks. Larger lease times will reduce the DHCPv4 traffic on the network and DHCPv4 server load.

6. Relaying DHCPv4 "over" IPv6

One of the broader contexts for preventing hosts from using IPv4 when IPv6 is sufficient is to only have IPv4 deployed on the network where necessary. Other parts of the network that don't require IPv4 will be native IPv6 only.

Should DHCPv4 relaying be used for these 255.255.255.255 subnet mask hosts, it would be expected that native IPv4 is used to carry DHCPv4 traffic between the DHCPv4 relay and server.

In keeping with the broader goal of minimising IPv4 deployment to just where it is necessary, an alternative and preferred method of relaying DHCPv4 to the DHCPv4 server would be over a point-to-point IPv4 in IPv6 tunnel [RFC2473], established directly between the DHCPv4 relay and the DHCPv4 server.

7. ISC Kea Configuration

Kea DHCPv4 server automatically derives the value of the Subnet Mask option, option value 1 from the ’subnet’ prefix option (e.g., the "/24" in 198.51.100.0/24). To cause Kea to supply a 255.255.255.255 subnet mask, it is necessary to use the ’libdhcp_flex_option.so’ hook library to modify the subnet mask supplied to clients. All other Kea configuration is conventional, with the exception of removing the configuration to supply clients with the Router Option, option value 3.

```
    "hooks-libraries": [
        {
            "library": "/usr/lib64/kea/hooks/libdhcp_flex_option.so",
                "parameters": {
                    "options": [
                        {
                            "code": 1,
                            "supersede": "255.255.255.255"
                        }
                    ]
                }
        }
    ],
```

8. Security Considerations

Although a host with a 255.255.255.255 subnet mask won't send traffic to other IPv4 hosts, excepting DHCPv4 traffic, it will still be able to receive packets from other hosts on the link if they are configured to do so. A 255.255.255.255 subnet mask host is still vulnerable to incoming unidirectional attacks. These hosts should have host based firewalling enabled on the interface that only allows DHCPv4 traffic between the DHCPv4 client and the DHCPv4 server or relay.

9. References

9.1. 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>.

9.2. Informative References

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

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

Acknowledgements

Review and comments were provided by Michael Richardson.

Author's Address

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