intarea
Internet-Draft
Intended status: Standards Track
Expires: 1 November 2026

R. S. van Mook
Asteroid International B.V.

30 April 2026

# IPv6-Resolved IPv4 Gateway
## draft-vanmook-intarea-ipv6-resolved-gateway-00

## Abstract

This document requests the allocation of the IPv4 special-purpose address 192.0.0.11/32 to enable IPv4 communication over IPv6-only networks without subnets, ARP, tunneling, or translation. Hosts configured with this address as their IPv4 default gateway resolve the next-hop link-layer address from the IPv6 neighbor cache instead of via ARP. IPv4 packets are carried natively, end-to-end.

## About This Document

This note is to be removed before publishing as an RFC.

Status information for this document may be found at https://datatracker.ietf.org/doc/draft-vanmook-intarea-ipv6-resolved-gateway/.

Source for this draft and an issue tracker can be found at https://github.com/remcovanmook/draft-ipv6-resolved-gateway.

## 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 1 November 2026.

Table of Contents



## 1. Introduction

Networks migrating to IPv6-only infrastructure still need to carry IPv4 traffic. Traditional mechanisms such as dual-stack, tunneling, and translation all reintroduce IPv4 at the infrastructure level. This document defines a special-purpose IPv4 address, 192.0.0.11, that signals to a host stack that link-layer resolution for the IPv4 default gateway is derived from the link-layer address entry for the IPv6 default router in the neighbor cache [RFC4861], rather than via ARP. The IPv4 routing table entry is unchanged; only the ARP resolution path is modified. This eliminates the need for IPv4 subnets and ARP on the local segment, removing the requirement for tunneling or translation at the first hop.

## 2. Conventions and Definitions

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.

## 3. Problem Statement

IPv4 next-hop resolution on a local link depends on ARP, which requires an IPv4 subnet to be configured on the link. In an IPv6-only network, no such subnet exists. Existing solutions -- dual-stack, tunneling, and translation-based approaches such as NAT64 -- generally require changes beyond the local segment. This document closes the first-hop resolution gap for IPv4 hosts on IPv6-only segments without requiring changes to host software, packet formats, or DHCPv4 clients.

## 4. Host Behavior and Next-Hop Resolution

When a host is configured to use 192.0.0.11 as its IPv4 default gateway, the host’s operating system MUST implement the following logic:

1. The host MUST maintain a functional IPv6 Neighbor Discovery implementation per [RFC4861] on the same interface, including default router discovery and neighbor cache maintenance. No additional action specific to this mechanism is required at interface configuration time.

2. When the next hop for an IPv4 packet is 192.0.0.11, the host MUST NOT perform ARP. Instead, it consults the IPv6 default router list and neighbor cache for the link-layer address, scoped to the interface on which 192.0.0.11 is configured as the IPv4 default gateway.

3. If the IPv6 default router link-layer address is in a usable NUD state (REACHABLE, STALE, DELAY, or PROBE per [RFC4861]), the IPv4 packet is sent in a link-layer frame addressed to that destination.

4. If no reachable IPv6 default router is known after the interface has completed initial configuration (i.e., at least one RA has been processed), the packet MAY be queued or dropped per implementation policy. If a last-known router address is available, a Neighbor Solicitation SHOULD be sent to that address. For behavior prior to first RA reception, see the startup paragraph below.

Host stacks MUST treat 192.0.0.11 as a reserved address requiring IPv6-based next-hop resolution, regardless of other address configuration on the interface. This behavior is unconditional and not dependent on any additional signaling. When a DHCPv4 lease configuring 192.0.0.11 expires and is not renewed, the host SHOULD remove 192.0.0.11 as the IPv4 default gateway and cease IPv6-based resolution on that interface. For statically configured deployments, removal is governed by local administrative policy.

Cross-interface resolution MUST NOT be performed. On multi-homed hosts, each interface independently resolves 192.0.0.11 against its own IPv6 neighbor cache state.

The following pseudocode defines the resolution logic:

```
on interface I (where 192.0.0.11 is the configured IPv4 gateway):

  if next-hop(pkt) == 192.0.0.11:

    routers = default_router_list(I)

    if routers is empty:
      if not first_ra_received(I):  /* startup: MUST queue */
        queue pkt
      else:                         /* mid-operation: MAY drop */
        queue or drop pkt
      send Router Solicitation on I  /* ff02::2; subject to
                                        RFC 4861 rate limiting */
      return

    selected = select from routers by:
      1. highest Default Router Preference (RFC 4191)
      2. NUD state == REACHABLE   /* if preference equal */
      3. implementation-defined   /* if reachability equal */

    lladdr = neighbor_cache(I, selected).lladdr

    if lladdr is valid and NUD state != INCOMPLETE:
      /* STALE, DELAY, and PROBE are usable; see RFC 4861 s7.3.3 */
      send pkt in link-layer frame with dst = lladdr
    else:
      queue or drop pkt           /* per implementation policy */
      send Neighbor Solicitation for selected router on I
```

Router selection uses Default Router Preference as defined in [RFC4191]. When multiple routers have equal preference and reachability, the tiebreaker is implementation-defined; use of most-recently-heard RA is one reasonable approach.

If the selected IPv6 default router becomes unreachable during an active session, the host SHOULD re-evaluate the default router list and select an alternative. Existing transport sessions will be disrupted if the link-layer next-hop changes; this is consistent with IPv6 router failure behavior and is not specific to this mechanism. Packets queued for a router that has become unreachable SHOULD be flushed and re-evaluated against the updated router selection.

Each time an interface transitions to an operational state and begins RA solicitation, a host may receive DHCPv4 configuration before any RA has been processed. Until a reachable IPv6 default router is known, IPv4 packets MUST be queued rather than silently dropped, pending RA reception. Implementations SHOULD bound this queue duration to avoid indefinite resource consumption. On queue timeout, packets SHOULD be dropped and an ICMPv4 Host Unreachable message MAY be generated toward the sending application.

# 5. Router Behavior

## 5.1. End-to-End Packet Flow

In this model, end hosts are assigned IPv4 addresses with a /32 prefix length. No IPv4 prefix is configured on the link; a sending host directs all IPv4 traffic to the first-hop router using the link-layer address derived from the IPv6 neighbor cache. A host cannot resolve another host's IPv4 address on the local link without router assistance; direct host-to-host IPv4 communication on the segment may occur via ICMPv4 redirect ([RFC1122], Section 3.2.2.2) from the gateway, but cannot be initiated by the host alone.

For return traffic to reach end hosts, operators MUST ensure that host /32 routes with an IPv6 next-hop per [RFC8950] are present in the routing infrastructure, allowing routers to forward IPv4 traffic toward the correct first-hop without requiring IPv4 addresses on any router interface. The mechanism by which the routing infrastructure learns these host routes is outside the scope of this document.

The following diagram illustrates the end-to-end packet flow. Router-to-router forwarding uses [RFC8950]; no IPv4 address is configured on any router interface. IPv4 addresses used in the diagram are from the documentation ranges defined in [RFC5737] and are not globally routable.

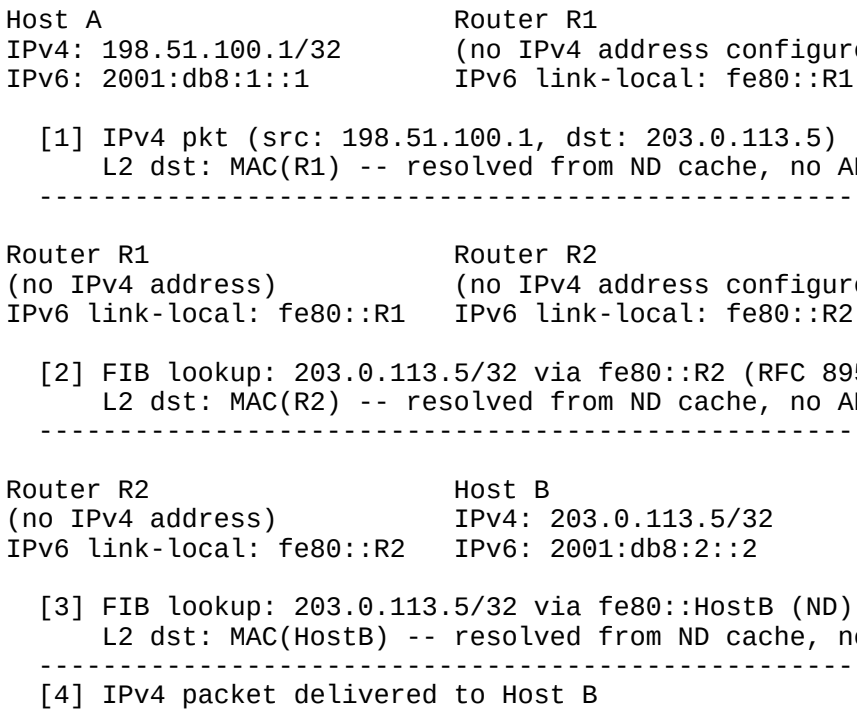


No ARP is exchanged at any point.

## 5.2. Router Ingress Behavior

Routers MUST treat 192.0.0.11 as an interface-scoped address -- valid only on the interface on which it is configured, and only for locally-terminated traffic. Specifically:

* It MUST NOT be injected into any routing protocol.

* It MUST NOT trigger overlapping-subnet checks.

* It MUST NOT appear as source or destination in any forwarded packet.

A router MAY respond to ICMPv4 echo requests addressed to 192.0.0.11 and MAY generate ICMPv4 Time Exceeded messages using 192.0.0.11 as the source address. All such messages are interface-local.

ICMPv4 error generation on IPv6-only transit routers is out of scope; see [RFC7600].

### 5.3. Backward Compatibility: Router ARP Response

The use of 192.0.0.11 as the DHCPv4 Router Option (Option 3) value is fully conformant with [RFC2132], which imposes no requirement that the router address be reachable via ARP on the same subnet.

Unmodified hosts receiving 192.0.0.11 as their IPv4 default gateway will issue an ARP request for it. A router SHOULD respond to such ARP requests with its own MAC address. This is not proxy ARP: no subnet exists, no remote host is being proxied.

This enables a two-tier deployment model on the same L2 segment:

* *Unmodified hosts:* router answers ARP; IPv4 forwarding works with zero host-side changes.

* *Updated hosts:* link-layer address resolved from IPv6 neighbor cache; ARP eliminated entirely.

Both tiers interoperate, allowing incremental deployment. The router requires no per-host state to support both tiers simultaneously: updated hosts will not send ARP requests for 192.0.0.11, so the router’s ARP response behavior is triggered only by unmodified hosts.

## 6. Deployment Considerations

This mechanism complements [RFC8925] (IPv6-Only Preferred Option). RFC 8925 allows hosts to signal a preference for IPv6-only operation, but operators must still provide IPv4 for hosts or applications that require it. This document provides exactly that fallback: IPv4 connectivity on an IPv6-only segment, without requiring a dual-stack infrastructure or an IPv4 subnet on the local link. A host that receives Option 108 and transitions to IPv6-only operation retains functional IPv4 connectivity via 192.0.0.11 without any additional configuration.

Hosts without an IPv6 stack are outside the scope of this document.

On segments with multiple routers advertising equal Default Router Preference -- common in datacenter ECMP fabrics -- hosts may make inconsistent router selections based on RA timing. Operators SHOULD configure explicit Default Router Preference values per [RFC4191] to ensure deterministic behavior.

Implementations in which IPv4 and IPv6 stacks are managed by separate processes -- as is common on mobile operating systems -- will require inter-process communication to expose the IPv6 neighbor cache to the IPv4 forwarding path. This is an implementation consideration and does not affect the on-wire behavior defined in this document.

## 7. Security Considerations

In the updated-host deployment model, ARP is eliminated from the network entirely. In the unmodified-host model, ARP is constrained to a single well-known address, reducing the ARP attack surface relative to conventional dual-stack networks.

The mechanism relies on the integrity of IPv6 Neighbor Discovery. Rogue RA risks apply as in any IPv6 deployment and can be mitigated with RA Guard [RFC6105]. Subnet scanning is mitigated since hosts carry /32 addresses only.

A host receiving 192.0.0.11 as its IPv4 default gateway on a network that does not implement this mechanism will issue an ARP request that receives no response, causing IPv4 connectivity to fail silently. Operators SHOULD ensure 192.0.0.11 is only offered via DHCPv4 on segments where the mechanism is deployed. DHCPv4 snooping and dynamic ARP inspection, where used, MUST be configured to permit ARP responses for 192.0.0.11 from the first-hop router.

This mechanism does not interact with IPv4 link-local address configuration per [RFC3927]. A host configured with 192.0.0.11 as its gateway and a link-local IPv4 source address will follow the same resolution logic defined in Section 4 (Host Behavior and Next-Hop Resolution).

## 8. Implementation Requirements

[PLACEHOLDER: This section will document conformance requirements and reference known implementations prior to IETF Last Call. An implementation is conformant if it satisfies all MUST and MUST NOT requirements in Sections 4 and 5. A reference implementation is available at https://github.com/remcovanmook/v4-with-v6-nh.]

## 9. IANA Considerations

This document requests that IANA assign 192.0.0.11/32 in the "IANA IPv4 Special-Purpose Address Registry" [RFC6890] as follows:

| Field | Value |
|---|---|
| Address Block | 192.0.0.11/32 |
| Name | IPv4 Gateway via IPv6 Resolution |
| RFC | This document |
| Allocation Date | (date of publication) |
| Termination Date | N/A |
| Source | True |
| Destination | True |
| Forwardable | False |
| Globally Reachable | False |
| Reserved-by-Protocol | False |

Table 1

The Destination=True designation reflects that 192.0.0.11 may appear as a destination in ICMPv4 messages received by the router on a local interface (see Section 5.2). It does not imply global reachability; Forwardable=False and Globally Reachable=False together preclude any use of this address beyond the local link.

## 10. References

### 10.1. 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/rfc/rfc2119>.

[RFC4191] Draves, R. and D. Thaler, "Default Router Preferences and More-Specific Routes", RFC 4191, DOI 10.17487/RFC4191, November 2005, <https://www.rfc-editor.org/rfc/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, <https://www.rfc-editor.org/rfc/rfc4861>.

[RFC6890] Cotton, M., Vegoda, L., Bonica, R., Ed., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, DOI 10.17487/RFC6890, April 2013, <https://www.rfc-editor.org/rfc/rfc6890>.

[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/rfc/rfc8174>.

[RFC8950] Litkowski, S., Agrawal, S., Ananthamurthy, K., and K. Patel, "Advertising IPv4 Network Layer Reachability Information (NLRI) with an IPv6 Next Hop", RFC 8950, DOI 10.17487/RFC8950, November 2020, <https://www.rfc-editor.org/rfc/rfc8950>.

## 10.2. Informative References

[RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, DOI 10.17487/RFC1122, October 1989, <https://www.rfc-editor.org/rfc/rfc1122>.

[RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, DOI 10.17487/RFC2132, March 1997, <https://www.rfc-editor.org/rfc/rfc2132>.

[RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", RFC 3927, DOI 10.17487/RFC3927, May 2005, <https://www.rfc-editor.org/rfc/rfc3927>.

[RFC5737] Arkko, J., Cotton, M., and L. Vegoda, "IPv4 Address Blocks Reserved for Documentation", RFC 5737, DOI 10.17487/RFC5737, January 2010, <https://www.rfc-editor.org/rfc/rfc5737>.

[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, <https://www.rfc-editor.org/rfc/rfc6105>.

[RFC7600] Despres, R., Jiang, S., Ed., Penno, R., Lee, Y., Chen, G., and M. Chen, "IPv4 Residual Deployment via IPv6 - A Stateless Solution (4rd)", RFC 7600, DOI 10.17487/RFC7600, July 2015, <https://www.rfc-editor.org/rfc/rfc7600>.

[RFC8925] Colitti, L., Linkova, J., Richardson, M., and T. Mrugalski, "IPv6-Only Preferred Option for DHCPv4", RFC 8925, DOI 10.17487/RFC8925, October 2020, <https://www.rfc-editor.org/rfc/rfc8925>.

## Appendix A. Acknowledgements

The author thanks Tobias Fiebig, Warren Kumari, and Jen Linkova.

## Author’s Address

Remco van Mook
Asteroid International B.V.
Email: remco@asteroidhq.com