

Internet Area Working Group
Internet-Draft
Intended status: Standards Track
Expires: 17 September 2026

J. Chroboczek
IRIF, Université Paris-Cité
W. Kumari
Google, LLC
T. Håland-Jørgensen
Red Hat
16 March 2026

IPv4 routes with an IPv6 next hop
draft-ietf-intarea-v4-via-v6-07

Abstract

V4-via-v6 routing is a technique that uses IPv6 next-hop addresses for routing IPv4 packets, and thus makes it possible to route IPv4 packets across a network where some routers have not been assigned IPv4 addresses. This document describes v4-via-v6 routing, and defines related operational procedures, notably the origination of ICMPv4 packets by nodes that might not have an IPv4 address.

About This Document

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

The latest revision of this draft can be found at <https://wkumari.github.io/draft-chroboczek-intarea-v4-via-v6/draft-ietf-intarea-v4-via-v6.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ietf-intarea-v4-via-v6/>.

Discussion of this document takes place on the Internet Area Working Group Working Group mailing list (<mailto:int-area@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/int-area/>. Subscribe at <https://www.ietf.org/mailman/listinfo/int-area/>.

Source for this draft and an issue tracker can be found at <https://github.com/wkumari/draft-chroboczek-intarea-v4-via-v6>.

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 17 September 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Conventions and Definitions	4
3. Operation	4
3.1. Structure of the routing table	5
3.2. Operation of the forwarding plane	5
3.3. Operation of routing protocols	5
4. ICMP Considerations	6
5. Implementation Status	7
5.1. Arista EOS	7
5.2. The Babel routing protocol	7
5.3. Linux	8
5.3.1. Example:	8
5.4. Mikrotik RouterOS	8
5.4.1. Example	9
5.5. Cisco NX-OS	9
6. Operational Considerations	9
7. Security Considerations	10
8. IANA Considerations	10
9. References	10
9.1. Normative References	10
9.2. Informative References	10
Acknowledgments	12
Changes	12
Version 06-07	12
Version 05-06	12

Version 04-05	13
Version 03-04	13
Version 02-03	13
Version 01-02	13
Version 00-01	13
Authors' Addresses	13

1. Introduction

The dominant form of routing in the Internet is next-hop routing, where a routing protocol constructs a routing table (or routing information base, RIB) which is used by a forwarding process to forward packets. The routing table is a data structure that maps network prefixes in a given family (IPv4 or IPv6) to next hops, pairs of an outgoing interface and a neighbor's network address, for example:

destination	next hop
2001:db8:0:1::/64	eth0, fe80::1234:5678
203.0.113.0/24	eth0, 192.0.2.1

When a packet is routed according to a given routing table entry, the forwarding plane typically maps the next-hop address to a link-layer address (a "MAC address") by using a neighbor discovery protocol (for example the Neighbor Discovery protocol (ND) [RFC4861] in the case of IPv6 over Ethernet, and the Address Resolution Protocol (ARP) [RFC0826] in the case of IPv4 over Ethernet). The link-layer address is then used to construct the link-layer frames that encapsulate forwarded packets.

It is apparent from the description above that there is no fundamental reason why the destination prefix and the next-hop address should be in the same address family: there is nothing preventing an IPv6 packet from being routed through a next hop with an IPv4 address (in which case the next hop's MAC address will be obtained using ARP), or, conversely, an IPv4 packet from being routed through a next hop with an IPv6 address. (In fact, it is even possible to store link-layer addresses directly in the next-hop entry of the routing table, thus avoiding the use of an address resolution protocol altogether, which was commonly done in networks using the OSI protocol suite.)

This document focuses on the specific case of routing IPv4 packets through an IPv6 next hop. This case is particularly interesting, since it makes it possible to build networks that have no IPv4 addresses except at the edges and still provide IPv4 connectivity to edge hosts. In addition, since an IPv6 next hop can use a link-local address that is autonomously configured, the use of such routes

enables a mode of operation where the network core has no statically assigned IP addresses of either family, which significantly reduces the amount of manual configuration required. (See also [RFC7404] for a discussion of the issues involved with such an approach.)

A route towards an IPv4 prefix that uses an IPv6 next hop is called a "v4-via-v6" route. V4-via-v6 routing is not restricted to routers, and could usefully be applied to hosts, but doing so would require solving the issue of host configuration, for example by extending either DHCPv4 or DHCPv6 to publish an IPv4 default route with an IPv6 next hop, which is out of scope for this document.

[RFC8950] discusses advertising of IPv4 Network Layer Reachability Information (NLRI) with a next-hop address that belongs to the IPv6 protocol, but confines itself to how this is carried and advertised in the BGP protocol. This document, on the other hand, discusses the concept of v4-via-v6 routes independently of any specific routing protocol, their design and operational considerations, and the implications of using them.

{ Editor note, to be removed before publication. This document is heavily based on draft-ietf-babel-v4viav6. When draft-ietf-babel-v4viav6 was going through IESG eval, Warren raised concerns that something this fundamental deserved to be documented in a separate, standalone document, so that it can be more fully discussed, and, more importantly, referenced cleanly in the future.}

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. Operation

Next-hop routing is implemented by two separate components, the routing protocol and the forwarding plane, that communicate through a shared data structure, the routing table.

3.1. Structure of the routing table

The routing table is a data structure that maps address prefixes to next hops, pairs of the form (interface, address). In traditional next-hop routing, the routing table maps IPv4 prefixes to IPv4 next hops, and IPv6 prefixes to IPv6 next hops. With v4-via-v6 routing, the routing table is extended so that an IPv4 prefix may map to either an IPv6 or an IPv4 next hop.

Resolution may be recursive: the next hop may itself be a prefix that requires further resolution to map to the outgoing interface and L2 address. V4-via-v6 routing does not prevent recursive resolution.

3.2. Operation of the forwarding plane

The forwarding plane is the part of the routing implementation that is executed for every forwarded packet. As a packet arrives, the forwarding plane consults the routing table, selects a single route matching the packet, and forwards the packet through the outgoing interface to the associated next-hop address.

With v4-via-v6 routing, the address family of the next-hop address is no longer determined by the address family of the prefix: since the routing table may map an IPv4 prefix to either an IPv4 or an IPv6 next hop, the forwarding plane must be able to determine, on a per-packet basis, which address resolution protocol (ARP for IPv4, ND for IPv6) to consult.

3.3. Operation of routing protocols

The routing protocol is the part of the routing implementation that is executed asynchronously from the forwarding plane, and whose role is to build the routing table. Since v4-via-v6 routing is a generalization of traditional next-hop routing, v4-via-v6 can interoperate with existing routing protocols: a traditional routing protocol produces a traditional next-hop routing table, which can be used by an implementation supporting v4-via-v6 routing.

However, in order to use the additional flexibility provided by v4-via-v6 routing, routing protocols need to be extended with the ability to populate the routing table with v4-via-v6 routes when an IPv4 address is not available or when the available IPv4 addresses are not suitable for use as a next hop.

Some protocols already support the advertisement of IPv4 routes with an IPv6 next hop, including Babel [RFC9229] and BGP [RFC8950]. Other protocols advertise both IPv4 and IPv6 prefixes over a single neighbor; these include:

- * Multi-Topology (MT) Routing in OSPF ([RFC4915])
- * Multi-Topology (MT) Routing in IS-IS ([RFC5120])

While both of these employ a common control plane, they use separate data planes, and therefore don't implement v4-via-v6 routing.

4. ICMP Considerations

The Internet Control Message Protocol (ICMPv4, or simply ICMP) [RFC0792] is a protocol related to IPv4 that is primarily used to carry diagnostic and debugging information. ICMPv4 packets may be originated by end hosts (e.g., the "destination unreachable, port unreachable" ICMPv4 packet), but they may also be originated by intermediate routers (e.g., most other kinds of "destination unreachable" packets).

Some protocols deployed in the Internet rely on ICMPv4 packets sent by intermediate routers. Most notably, path MTU Discovery (PMTUD) [RFC1191] is an algorithm executed by end hosts to discover the maximum packet size that a route is able to carry. While there exist variants of PMTUD that are purely end-to-end [RFC4821], the variant most commonly deployed in the Internet has a hard dependency on ICMPv4 packets originated by intermediate routers: if intermediate routers are unable to send ICMPv4 packets, PMTUD may lead to persistent black-holing of IPv4 traffic.

A router must therefore be able to generate ICMP Destination Unreachable messages ([RFC1812] Section 5.2.7.1). The source address of these messages must be one of the addresses assigned to the outgoing interface; if no such address has been assigned, then one of the other addresses assigned to the router, known as the "router-id", must be used ([RFC1812] Section 4.3.2.4).

Routers implementing the mechanism described in this document do not need to have any IPv4 addresses assigned to any of their interfaces, and [RFC1812] does not specify what happens if no router-id has been assigned. If a router does not have any IPv4 addresses assigned, the router MUST use the dummy address 192.0.0.8 as the source address of outgoing ICMP packets ([RFC7600], Section 4.8, Requirement R-22).

Using the dummy address as the source of ICMPv4 packet causes a number of drawbacks:

- * using the same address on multiple routers may hamper debugging and fault isolation, e.g., when using the `_traceroute_` utility (but see [I-D.draft-ietf-intarea-extended-icmp-nodeid] for a possible solution to this problem);

- * packets originating from 192.0.0.8 might be considered as spoofed traffic and dropped by firewalls at network boundaries.

For these reasons, even if a router performs v4-via-v6 routing on all interfaces, it MAY be assigned one or more IPv4 addresses.

5. Implementation Status

(RFC Editor: please remove this section before publication.)

(This section records the status of known implementations of the protocol defined by this specification at the time of writing, and is based on a proposal described in RFC 7942. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.)

5.1. Arista EOS

Arista has supported static IPv4 routes with IPv6 next hops since EOS-4.30.1.

5.2. The Babel routing protocol

As noted above, this document is heavily based on RFC9229 (nee draft-ietf-babel-v4viav6), and this functionality is supported by babeld.

Pasted below is email sent to the babel mailing list (archived at <https://mailarchive.ietf.org/arch/msg/babel/QtFi3F4TFfF7fXXlkHSpEnuT44Y/>)

An IPv4 route across three nodes with IPv6 addresses only:

```
$ ip route show 10.0.0.2
10.0.0.2 via inet6 fe80::216:3eff:fe00:1 dev lxcbr0 proto babel onlink
```

Here's how it's logged by babeld:

```
10.0.0.2/32 from 0.0.0.0/0 metric 384 (384) refmetric 288 id
02:16:3e:ff:fe:9a:5e:22 seqno 36425 chan (255) age 15 via lxcbr0 neigh
fe80::216:3eff:fe00:1 (installed)
```

Traceroute is a little confusing:

```
$ traceroute 10.0.0.2
traceroute to 10.0.0.2 (10.0.0.2), 30 hops max, 60 byte packets
 1  192.0.0.8 (192.0.0.8)  0.079 ms  0.019 ms  0.014 ms
 2  192.0.0.8 (192.0.0.8)  0.040 ms  0.023 ms  0.042 ms
 3  192.0.0.8 (192.0.0.8)  0.061 ms  0.030 ms  0.030 ms
 4  10.0.0.2 (10.0.0.2)  0.060 ms  0.040 ms  0.039 ms
```

PMTUD works fine (thanks to Toke):

```
19:58:47.402871 IP 192.168.0.27.60046 > 10.0.0.2.22: Flags [.] \
seq 33:1481, ack 33, win 502, options [nop,nop,TS val 917354570 \
ecr 1849974691], length 1448
19:58:47.402874 IP 192.168.0.27.60046 > 10.0.0.2.22: Flags [P.] \
seq 1481:1537, ack 33, win 502, options [nop,nop,TS val 917354570 \
ecr 1849974691], length 56
19:58:47.402906 IP 192.0.0.8 > 192.168.0.27: ICMP 10.0.0.2 \
unreachable- need to frag (mtu 1420), length 556
19:58:47.402919 IP 10.0.0.2.22 > 192.168.0.27.60046: Flags [.] \
ack 33, win 509, options [nop,nop,TS val 1849974692 \
ecr 917354569,nop,nop,sac 1 {1481:1537}], length 0
19:58:47.402934 IP 192.168.0.27.60046 > 10.0.0.2.22: Flags [.] \
seq 33:1401, ack 33, win 502, options [nop,nop,TS val 917354570 \
ecr 1849974692], length 1368
```

-- Juliusz

5.3. Linux

Linux has supported v4-via-v6 routes since kernel version 5.2, released on 2019-07-07.

5.3.1. Example:

```
rincewind ~ #
ip -4 r a 192.0.2.23/32 via inet6 2001:db8::2342

rincewind ~ # ip r s 192.0.2.23/32
192.0.2.23 via inet6 2001:db8::2342 dev wlp36s0.25
```

5.4. Mikrotik RouterOS

Mikrotik RouterOS has supported v4-via-v6 routes since (at least) version 7.11beta2

{Editor note: I'm not sure when support was added. I tested this in Version 7.11beta2, and it worked there, but I believe that this functionality has existed for a while. I'll try to find out when it was added.}

5.4.1. Example

```
[wkumari@Dulles-CCR] /ip/route> print
Flags: D - DYNAMIC; I - INACTIVE, A - ACTIVE; c - CONNECT, s - STATIC,
d -DHCP, v - VPN; H - HW-OFFLOADED
Columns: DST-ADDRESS, GATEWAY, DISTANCE
#      DST-ADDRESS      GATEWAY      DISTANCE
0  As  192.0.2.0/24      fe80::201:5cff:feb2:1646%1_Comcast  1
```

5.5. Cisco NX-OS

Cisco NX-OS has supported v4-via-v6 routes "for more than 8 years" --
Krishnaswamy Ananthamurthy

6. Operational Considerations

V4-via-v6 routing makes it easy to route IPv4 traffic across interfaces that have not been assigned IPv4 addresses, and therefore has the potential to reduce the number of IPv4 addresses consumed and hopefully simplify the management of double-stack networks. Since it promises IPv4 routing essentially "for free" once IPv6 addressing has been set up, it has the potential to slightly accelerate the deployment of IPv6.

Just like any other extension to an existing technology, however, it requires changes to existing infrastructure. Even though v4-via-v6 routes are similar in structure to traditional next-hop routes, at least some monitoring and management tools will not be able to interpret them. Deployment of v4-via-v6 routing in a network requires testing and potentially updating of all tools and scripts that manipulate or examine routes.

V4-via-v6 routing encourages a model of deployment where some routers have no IPv4 addresses even though they forward IPv4 traffic. Such routers make debugging of IPv4 routing issues somewhat more difficult, most notably by making the output of the `_traceroute_` utility less informative than it would otherwise be (see Section 4). Even if the procedures described in [I-D.draft-ietf-intarea-extended-icmp-nodeid] are deployed on all such routers, older versions of `_traceroute_` will not be able to interpret the additional information. Network administrators might want to provision IPv4 addresses on all routers in order to simplify debugging.

7. Security Considerations

The techniques described in this document make routing more flexible by allowing IPv4 routes to propagate across a section of a network that has only been assigned IPv6 addresses. This additional flexibility might invalidate otherwise reasonable assumptions made by network administrators, which could potentially cause security issues.

For example, if an island of IPv4-only hosts is separated from the IPv4 Internet by routers that have not been assigned IPv4 addresses, a network administrator might reasonably assume that the IPv4-only hosts are unreachable from the IPv4 Internet. This assumption is broken if the intermediary routers implement v4-via-v6 routing, which might make the IPv4-only hosts reachable from the IPv4 Internet. If this is not desirable, then the network administrator must filter out the undesirable traffic in the forwarding plane by implementing suitable packet filtering rules.

8. IANA Considerations

No IANA actions are requested by this document.

9. References

9.1. Normative References

- [RFC1812] Baker, F., Ed., "Requirements for IP Version 4 Routers", RFC 1812, DOI 10.17487/RFC1812, June 1995, <<https://www.rfc-editor.org/rfc/rfc1812>>.
- [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>>.
- [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>>.

9.2. Informative References

- [I-D.draft-ietf-intarea-extended-icmp-nodeid]
Fenner, B. and R. Thomas, "Adding Extensions to ICMP Errors for Originating Node Identification", Work in Progress, Internet-Draft, draft-ietf-intarea-extended-icmp-nodeid-04, 19 August 2025,
<<https://datatracker.ietf.org/doc/html/draft-ietf-intarea-extended-icmp-nodeid-04>>.
- [RFC0792] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, DOI 10.17487/RFC0792, September 1981,
<<https://www.rfc-editor.org/rfc/rfc792>>.
- [RFC0826] Plummer, D., "An Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, DOI 10.17487/RFC0826, November 1982,
<<https://www.rfc-editor.org/rfc/rfc826>>.
- [RFC1191] Mogul, J. and S. Deering, "Path MTU discovery", RFC 1191, DOI 10.17487/RFC1191, November 1990,
<<https://www.rfc-editor.org/rfc/rfc1191>>.
- [RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, DOI 10.17487/RFC4821, March 2007,
<<https://www.rfc-editor.org/rfc/rfc4821>>.
- [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>>.
- [RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", RFC 4915, DOI 10.17487/RFC4915, June 2007,
<<https://www.rfc-editor.org/rfc/rfc4915>>.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, DOI 10.17487/RFC5120, February 2008,
<<https://www.rfc-editor.org/rfc/rfc5120>>.
- [RFC7404] Behringer, M. and E. Vyncke, "Using Only Link-Local Addressing inside an IPv6 Network", RFC 7404, DOI 10.17487/RFC7404, November 2014,
<<https://www.rfc-editor.org/rfc/rfc7404>>.

- [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>>.
- [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>>.
- [RFC9229] Chroboczek, J., "IPv4 Routes with an IPv6 Next Hop in the Babel Routing Protocol", RFC 9229, DOI 10.17487/RFC9229, May 2022, <<https://www.rfc-editor.org/rfc/rfc9229>>.

Acknowledgments

This document is based on [RFC9229], which was produced by the IETF Babel working group.

We are grateful to Joe Abley, Krishnaswamy Ananthamurthy, Vint Cerf, Joe Clarke, Lorenzo Colitti, Bill Fenner, Tobias Fiebig, John Gilmore, Bob Hinden, Jen Linkova, David Lamparter, Gyan Mishra, tom patch, Herbie Robinson, Behcet Sarikaya, David Schinazi, Ole Troan, and テ詠ic Vyncke for helpful comments and suggestions about this document.

Changes

This section is to be removed before publication, and the primary change log is the git repository. This is just a place to note some of the more substantive changes.

Version 06-07

- * Made the reference to RFC 7600 informative.
- * Downgraded the suggestion to assign an IPv4 address from SHOULD to MAY.
- * Changed the first author's affiliation (but don't get used to it, they'll likely rename us again).

Version 05-06

- * Rewrote abstract.
- * Added RFC 7942 boilerplate.

- * Added positive factors to Operational Considerations.
- * Editorial changes.

Version 04-05

- * Minor editorial changes.

Version 03-04

- * Added a section about operational considerations.
- * Made it clear that ARP/ND are not necessarily used.
- * Removed any mention of v4-only, since it's not quite correct that v4-via-v6 is v4-only.

Version 02-03

- * Warren is a smart guy, but he still pushed a branch instead of the main one, so -03 is actually what -02 should have been.

Version 01-02

- * Addressed comments from Vint and Jen.

Version 00-01

- * Added note that this works just as well for IPv6 routes with an IPv4 next hop. (テ詠ic Vyncke)
- * Cisco NX-OS has supported v4-via-v6 routes "for more than 8 years" (Krishnaswamy Ananthamurthy)
- * Mention recursive next hops, and that the next hop may be a prefix. (Krishnaswamy Ananthamurthy)
- * Hosts are routers too! (David Lamparter)
- * Removed the claim that it's mainly a UI issue.

Authors' Addresses

Juliusz Chroboczek
IRIF, Universitテウ Paris-Citテウ
Case 7014
75205 Paris Cedex 13
France

Email: jch@irif.fr

Warren Kumari
Google, LLC
Email: warren@kumari.net

Toke Hテ ク iland-Jテ ク rgensen
Red Hat
Email: toke@toke.dk