Announcing IPv4 routes with an IPv6 next-hop in the Babel routing protocol

aka. draft-bastian-babel-v4ov6

Théophile Bastian, joint work with Juliusz Chroboczek

ENS Paris, IRIF, Nexedi

Traditional routing

Babel: routing protocol. Only concern: build the routing table.

Routing table

Network Prefix (IPv6)	→	Next-Hop (IPv6)
2001::/64		fd80::42

Network Prefix (IPv4)	→	Next-Hop (IPv4)
10.0.0.0/24		10.0.0.1
10.0.0.0/16		10.0.1.1
10.0.0.0/8		~~fe80::f0~~

The router's job

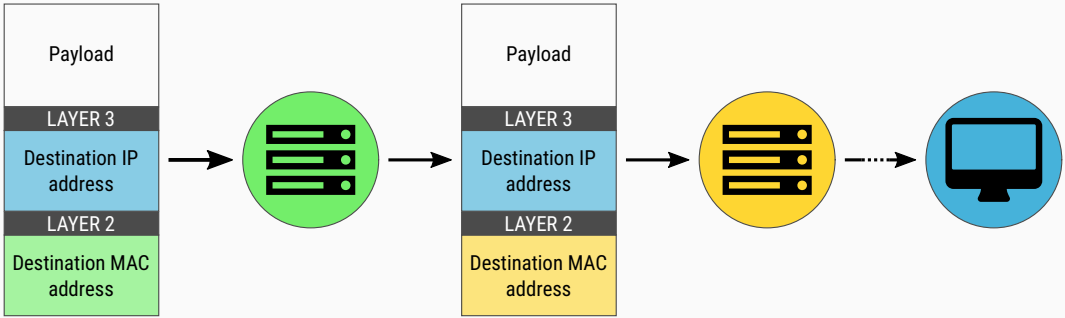

What's under the hood?

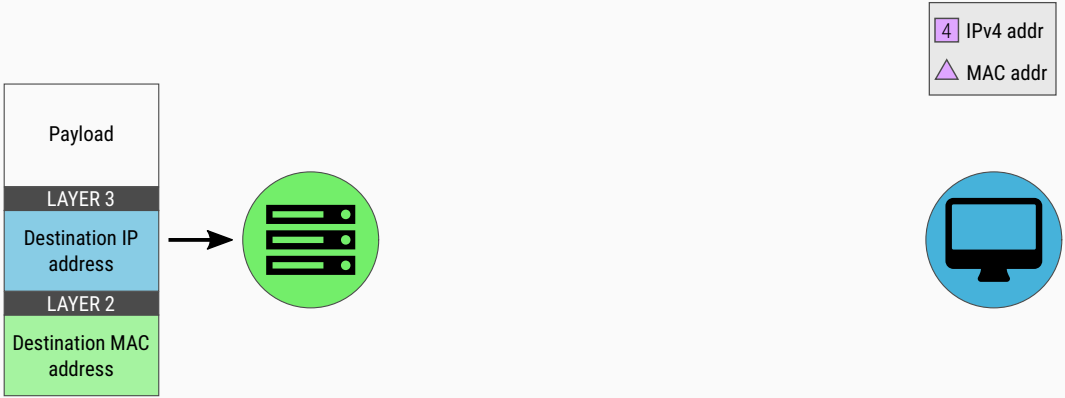

What's under the hood?

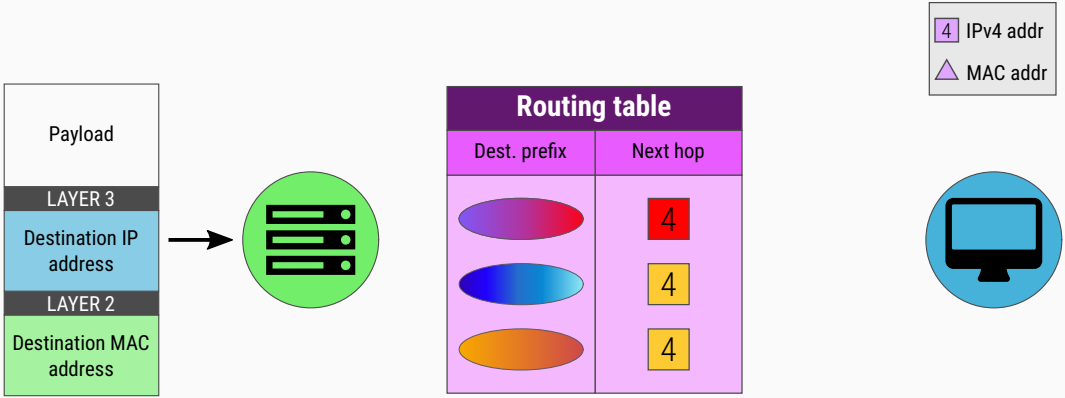

What's under the hood?

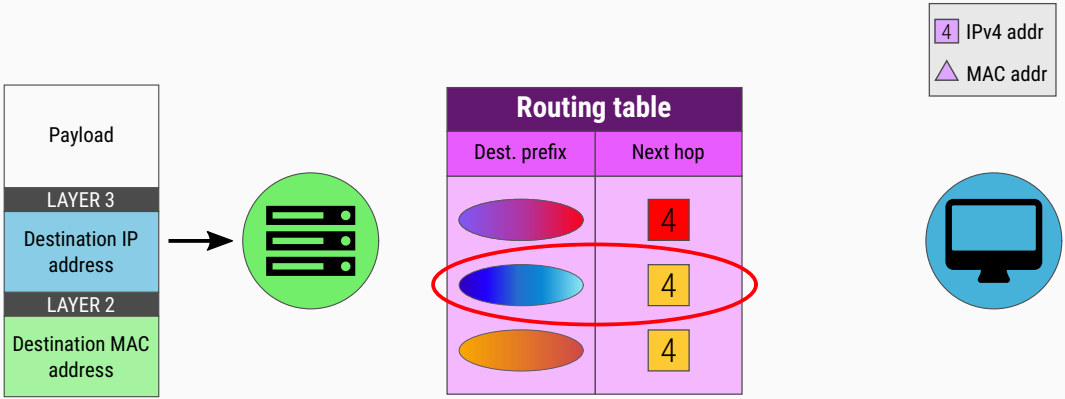

What's under the hood?

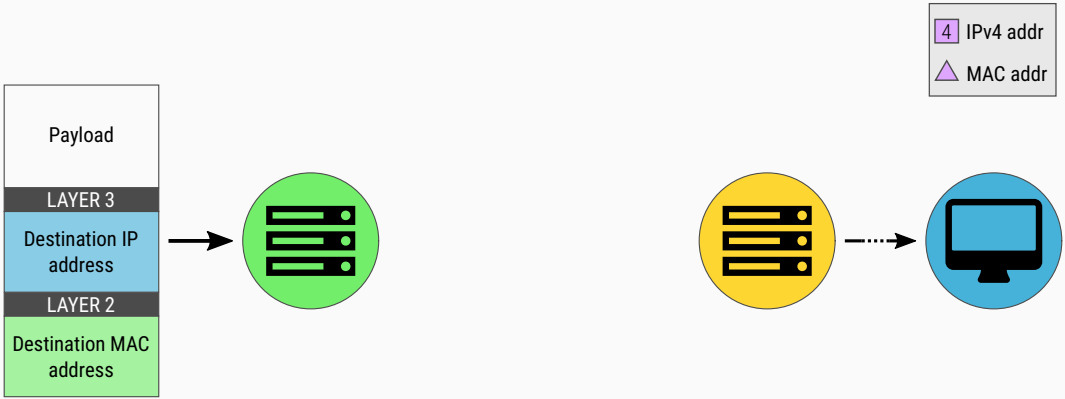

What's under the hood?

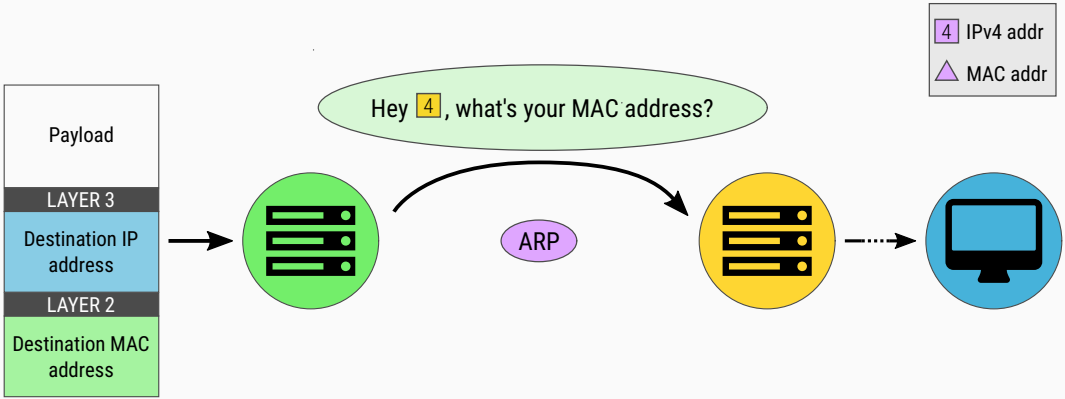

What's under the hood?

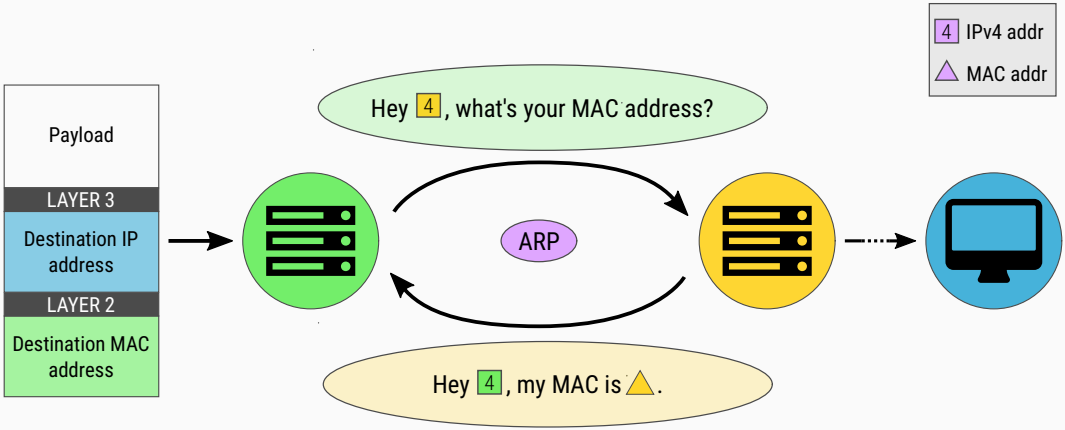

What's under the hood?

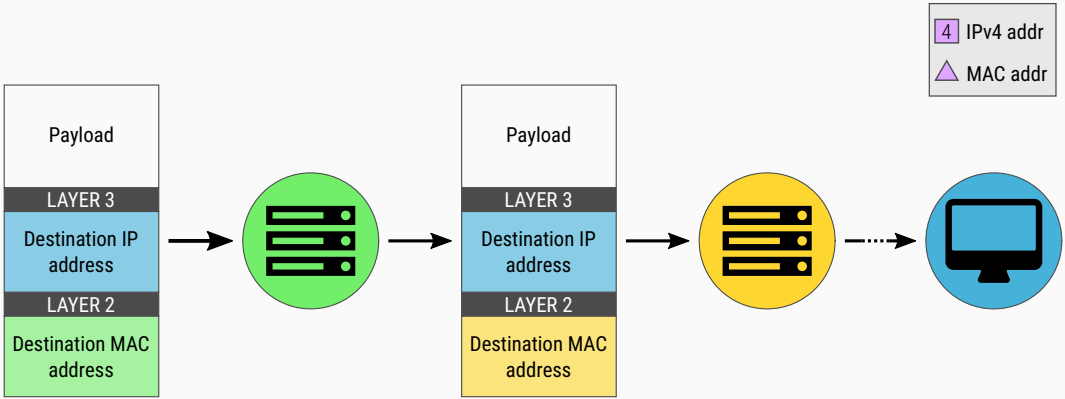

What's under the hood?

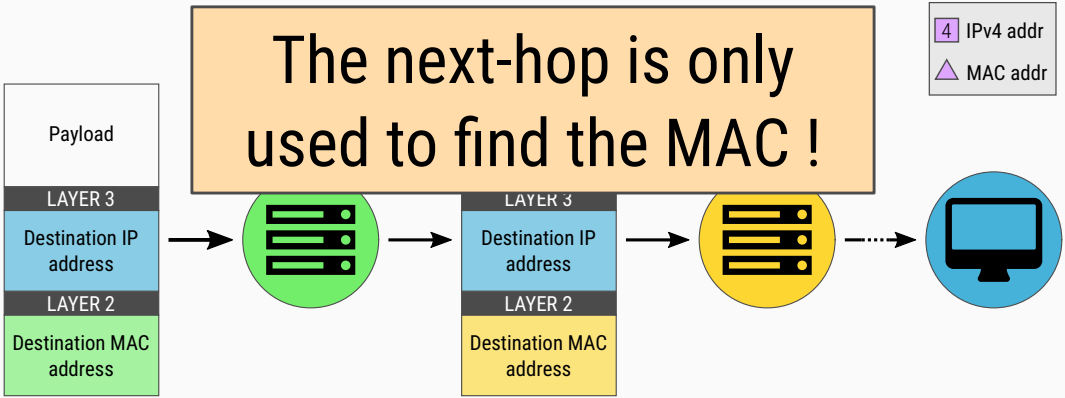

What's under the hood?

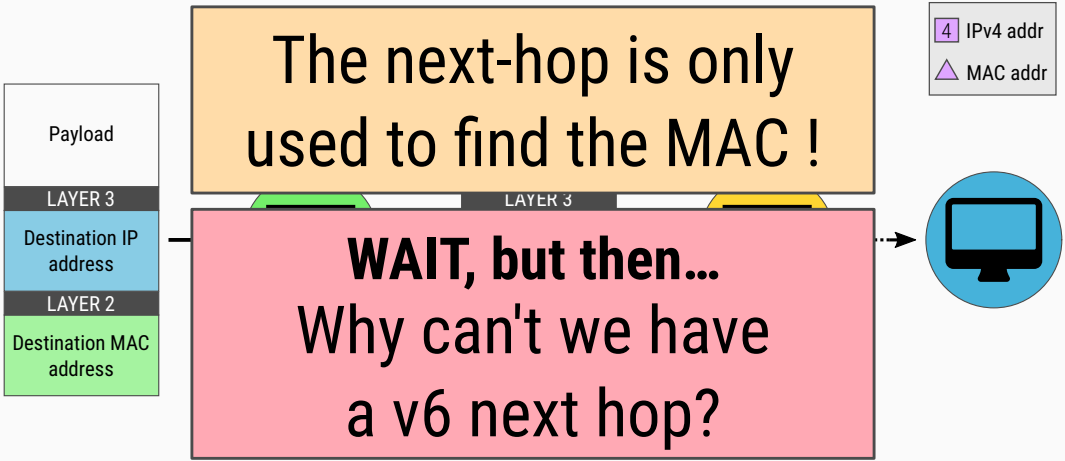

What's under the hood?

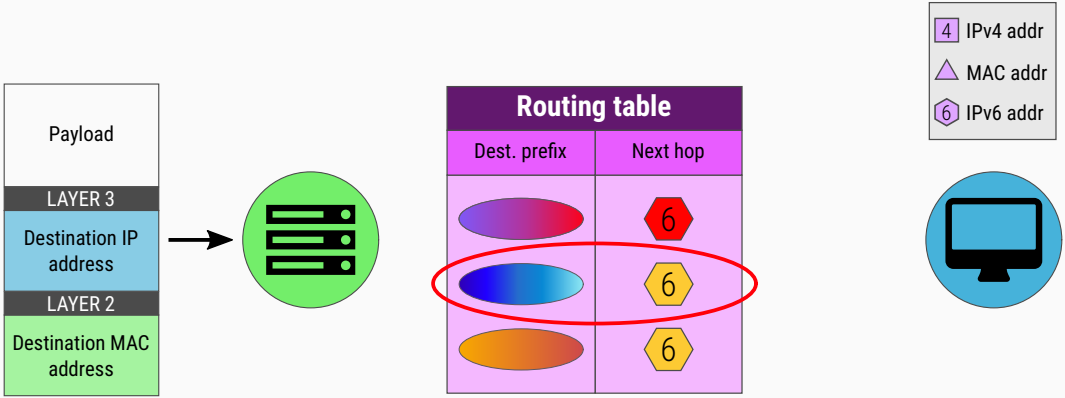

What's under the hood?

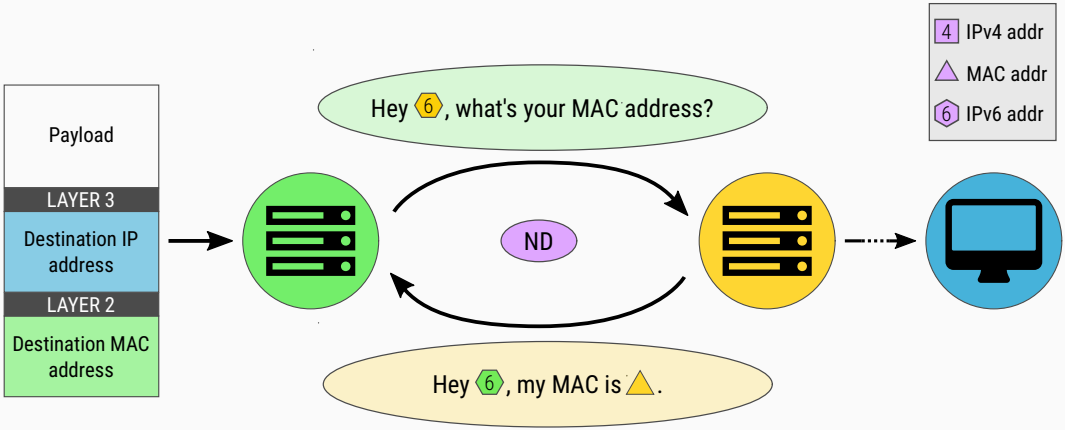

What's under the hood?

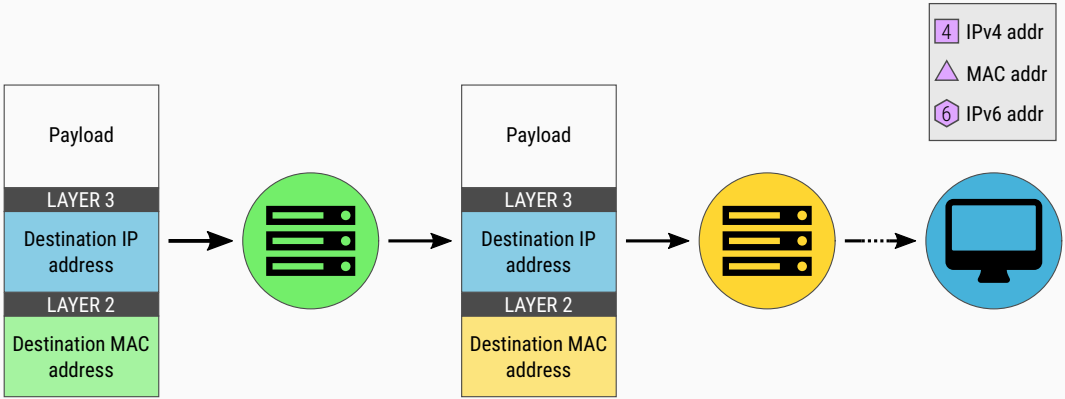

The idea behind

New type of route! v4 prefix with v6 next-hop: v4-over-v6 route.

- Useful to go through a router with no v4 address configured
- Possible use: v6-only core, still serve v4 to clients *without tunnel*

Not an original idea! BGP had it first: draft-ietf-bess-rfc5549revision

Linux support

Supported in Linux since 5.2 – July 2019!

git **d1566268 – 2019-04-05 – ipv4: Allow ipv6 gateway with ipv4 routes**
Add support for `RTA_VIA` and allow an IPv6 nexthop for v4 routes.

- Usually, add a v4 route with

```
# ip route add 10.42.0.0/16 via 10.40.0.42 dev eno1
```

- v4-over-v6 (recent kernel):

```
# ip route add 10.42.0.0/16 via inet6 fe80::a0de:baf:b39b dev eno1
```

Babel protocol extension

Advertising v4-over-v6 routes

Interface with v4 address: no changes; just as unextended babel.

Interface with only v6 addresses:

- Receiving a v4 route: install it anyway, the router's address is irrelevant.
- Announcing a v4 route: use v4-over-v6; we need a next-hop and only have v6 at hand.

Backwards compatibility

Backwards compatibility: an unextended node must

- Be able to ignore v4-over-v6 routes
- Route correctly pure v4 and v6

In *unextended* Babel, Address Encodings (AEs) define the type of address/prefix contained in a TLV.

- IPv4 address
- IPv6 address
- link-local IPv6 address

Encoding v4-over-v6 routes: choices

Various encodings possible, among which:

1. *[Toke]* An IPv4 route announced without previously setting a valid next-hop is considered v4-over-v6.
   - Is backwards compatible – but not obviously so
   - Not clear an extension is being used
2. *[Bastian + JCH]* New specific AE for v4-over-v6 routes, and next-hop is...
   - i ...specific to v4-over-v6 routes
   - ii ...same as for v6 routes
     - → more compact, just as clear
   - Backwards compatible and clean
   - A tiny bit more verbose

Encoding v4-over-v6 routes: choices

Various encodings possible, among which:

1. *[Toke]* An IPv4 route announced without previously setting a valid next-hop is considered v4-over-v6.
   - Is backwards compatible – but not obviously so
   - Not clear an extension is being used
2. *[Bastian + JCH]* New specific AE for v4-over-v6 routes, and next-hop is...
   - i ...specific to v4-over-v6 routes
   - ii **...same as for v6 routes**
     - → more compact, just as clear
   - Backwards compatible and clean
   - A tiny bit more verbose

⇝ **2.ii. Add a new AE. No need for new TLVs.**

Conclusion

- New type of route: v4-over-v6, v4 destination, v6 next-hop
- Route IPv4 over an IPv6 network core. *Look, Ma! No tunnels!*
- Protocol described and drafted
- Production-ready implementation available on the babeld repository

- **Intended status: experimental**
- Opinions: should it be adopted by workgroup or carried alone?

RFC draft

`huit.re/draft-v4ov6`

These slides

`huit.re/ietf108-v4ov6`