draft-equinox-intarea-dhcpv4-route4via6

DHCPv4 Option for IPv4 routes with IPv6 nexthops

David ‘equinox’ Lamparter

NetDEF, Inc.

Tobias Fiebig

Max-Planck Institut für Informatik

Problem Statement

- We can route IPv4 traffic via IPv6 nexthops (see also: draft-ietf-intarea-v4-via-v6)
- This allows clean IPv6-first network architectures with IPv4 as an add-on

Problem: How can we configure clients that do need IPv4?

What we want: An IPv4-less(ish) Network

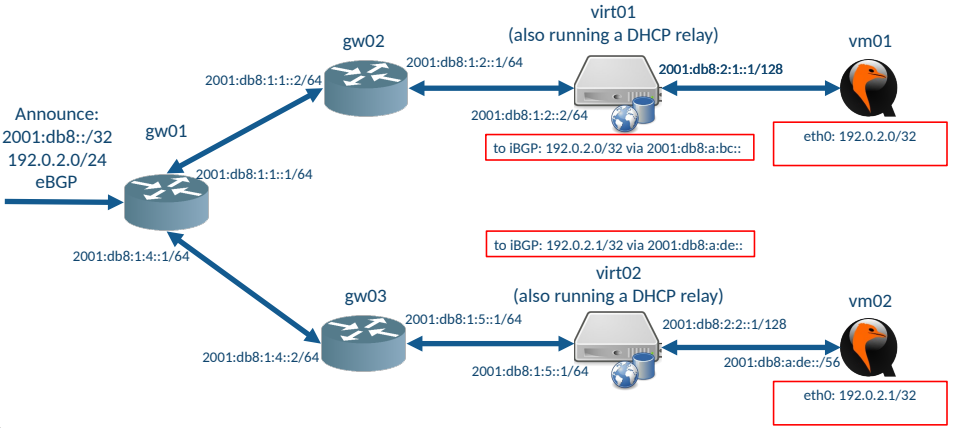

Solution Step 1: RFC7341

RFC7341: DHCPv4-over-DHCPv6 (DHCP 4o6) Transport

- Allows doing DHCPv4 tunneled over DHCPv6
- Removes the need for having v4 addresses for a DHCPv4 relay/server

Remaining Problem: Communicating v4v6 (default) route(s) to clients

DHCPv4-over-DHCPv6 Packet

```
▸ Internet Protocol Version 6, Src: fe80::f066:7cff:fe70:9e32, Dst: ff02::1:2
▸ User Datagram Protocol, Src Port: 546, Dst Port: 547
▾ DHCPv6
    Message type: 4o6 Query (20)
    Transaction ID: 0x000000
  ▾ DHCPv4 Message
      Option: DHCPv4 Message (87)
      Length: 300
    ▾ Dynamic Host Configuration Protocol (Request)
        Message type: Boot Request (1)
        Hardware type: Ethernet (0x01)
        Hardware address length: 6
        Hops: 0
        Transaction ID: 0x5cd92337
        Seconds elapsed: 0
      ▸ Bootp flags: 0x0000 (Unicast)
        Client IP address: 0.0.0.0
        Your (client) IP address: 0.0.0.0
        Next server IP address: 0.0.0.0
        Relay agent IP address: 0.0.0.0
        Client MAC address: f2:66:7c:70:9e:32 (f2:66:7c:70:9e:32)
```

(Default) Routes in an IPv4-less(ish) Network

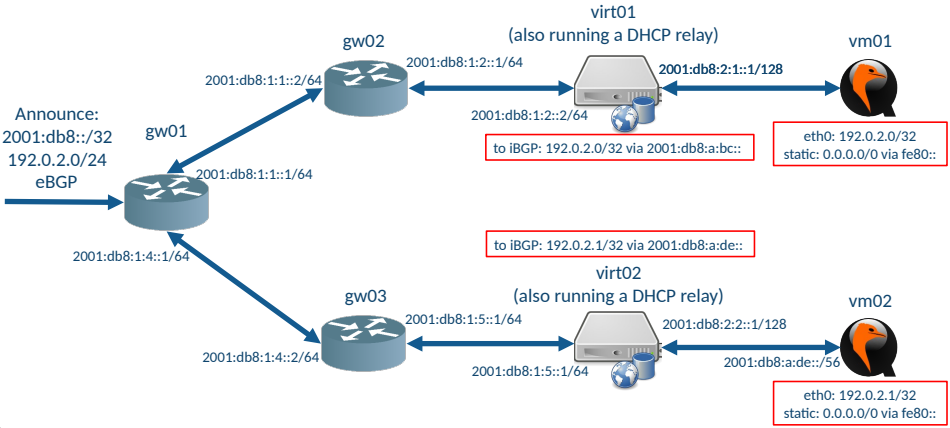

Solution Step 2: DHCPv4 Option for IPv4 Routes with IPv6 Nexthops

Adding a new DHCP**v4** (Yes, **v4** and **not** v6) option to communicate v4 routes with v6 nexthops to clients

- Allow multiple routes
- New option with multiple nexthop pairs

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Type (TBA1)  |     Length    |     Prefix/nexthop pairs      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+                               |
:                                  ...                          :
```

Solution Step 2: DHCPv4 Option for IPv4 Routes with IPv6 Nexthops

- Type: TBA1 (field defined in [DHCP-OPT])
- Length: as defined in [DHCP-OPT]
- Prefix/nexthop pairs: zero or more items as defined below. Note multiple options are concatenated before processing.

Solution Step 2: Destination Prefix and Nexthop Pair

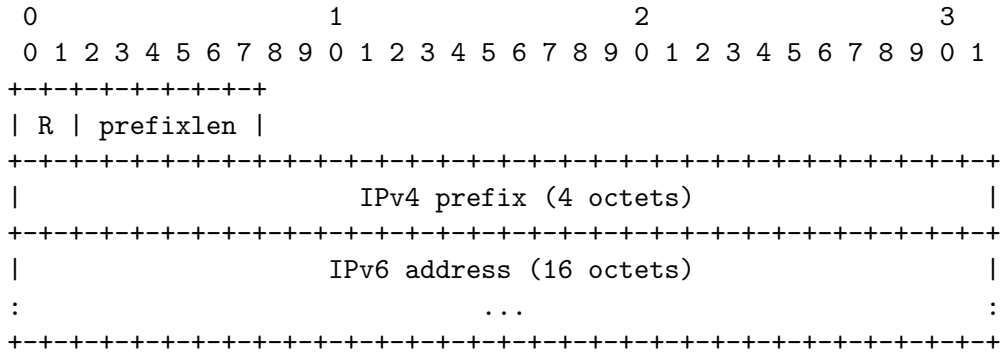

Solution Step 2: Destination Prefix and Nexthop Pair

- R: Reserved bits. MUST be sent as zero, MUST be ignored on receipt.
- prefixlen: IPv4 prefix length, integer value from 0 to 32 (inclusive)
- IPv4 prefix: The route's destination prefix, with unused bits zeroed. Validity and behavior for specific values is described in Section 3.2.
- IPv6 address: The IPv6 address specifying the nexthop for this route. Refer to Section 3.1 for valid values and associated behavior.

Next Steps

- Anyone interested in working on this?
- Call for adoption?
- Anyone up for implementing this?