draft-perreault-sunset4-noipv4-00

Simon Perreault <simon.perreault@viagenie.ca>
Wes George <wesley.george@twcable.com>
Tina Tsou <tina.tsou.zouting@huawei.com>

IETF 84, Vancouver
SUNSET4 meeting

2012-07-30

Introduction

- Dual-stack node
- No DHCPv4 response == failure condition
  - Retransmit endlessly
  - Some implementations reject the connection entirely, including IPv6

Home Router Example

- Boot process:
  - Assign IPv4 address on LAN interface
  - Start LAN services (e.g., DNS)
  - Starts handing out IPv4 addresses on LAN
- This is done unconditionally, without taking the status of the IPv4 connectivity on the WAN interface into account.
- Hosts on the LAN:
  - Install a default route pointing to the router
  - Start behaving as if IPv4 connectivity was available
  - IPv4 packets to the Internet get dropped at the router
  - Timeouts happen
- End result: IPv4 is fully active on the LAN and on the router itself even when it is desired that it be turned off.

There is a need

- A new mechanism to indicate the absence of IPv4 connectivity
- Must be transported over IPv6
  - ...since the goal is eliminating IPv4

The No-IPv4 DHCPv6 option

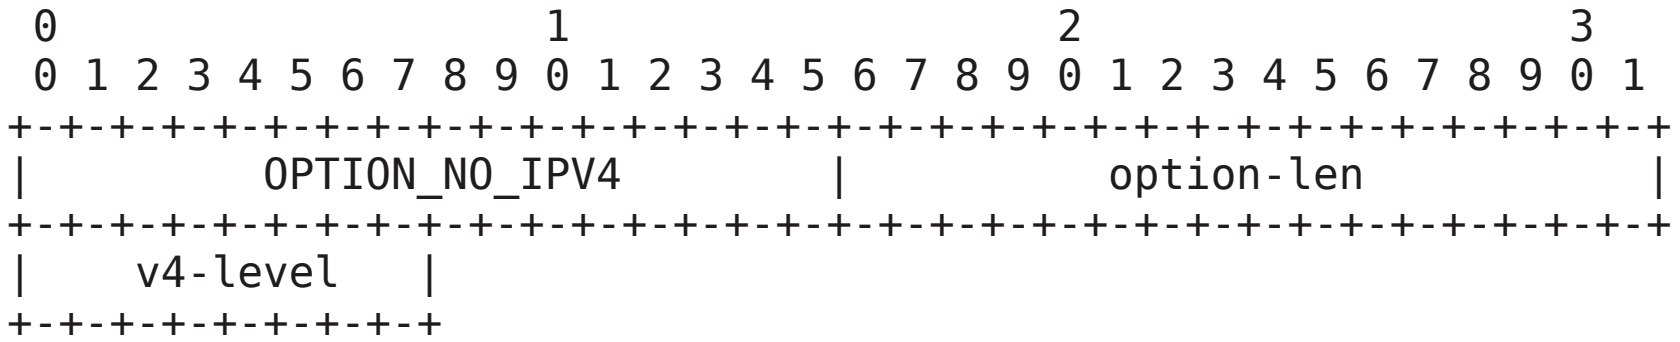

option-code OPTION_NO_IPV4 (TBD).
option-len 1.
v4-level Level of IPv4 functionality.

Terminology

- Upstream Interface: An interface on which the No-IPv4 DHCPv6 option is received by a DHCPv6 client.

v4-level

- 0 – IPv4 fully enabled
  - Equivalent to the absence of the No-IPv4 option.
  - Purpose: so that a DHCPv6 server can explicitly re-enable IPv4 access by including it in a Reply message following a Reconfigure.

v4-level

- 1 – No IPv4 upstream, local IPv4 permitted
  - IPv4 connectivity is unavailable on the link on which this option is received.
  - The following IPv4 functionality MUST be disabled on the upstream interface:
    A. IPv4 addresses MUST NOT be assigned.
    B. Currently-assigned IPv4 addresses MUST be unassigned.
    C. Dynamic configuration of link-local IPv4 addresses [RFC3927] MUST be disabled.
    D. IPv4, ICMPv4, or ARP packets MUST NOT be sent.
    E. IPv4, ICMPv4, or ARP packets received MUST be ignored.
    F. DNS A queries MUST NOT be sent, even transported over IPv6.

v4-level

- 1 – continued...
  - If all DHCPv6-configured interfaces receive the No-IPv4 option with value 1 or 2, and no other interface provides IPv4 connectivity to the Internet, IPv4 is ***partially*** shut down.
    - Upstream interface: see previous slide
    - Other interfaces: IPv4 addresses MUST NOT be assigned except for the following ranges:
      - Loopback (127.0.0.0/8)
      - Link-local (169.254.0.0/16)
      - Private-use (RFC 1918)

v4-level

- 2 – No IPv4 at all
  - Disable IPv4 functionality on upstream interface as for level 1.
  - If all DHCPv6-configured interfaces received the No-IPv4 option with exclusively value 2, and no other interface provides IPv4 connectivity to the Internet, IPv4 is completely shut down.
    A. IPv4 address MUST NOT be assigned to any interface.
    B. Currently-assigned IPv4 addresses MUST be unassigned.
    C. Dynamic configuration of link-local IPv4 addresses [RFC3927] MUST be disabled.
    D. IPv4, ICMPv4, or ARP packets MUST NOT be sent on any interface.
    E. IPv4, ICMPv4, or ARP packets received on any interface MUST be ignored.
    F. In the above, "any interface" includes loopback interfaces. In particular, the 127.0.0.1 special address MUST be removed.
    G. Server programs listening on IPv4 addresses (e.g., a DHCPv4 server) MAY be shut down.
    H. DNS A queries MUST NOT be sent, even transported over IPv6.
    I. If the host or router also runs a DHCPv6 server, it SHOULD include the No-IPv4 option with value 2 in DHCPv6 responses it sends to clients that request it, unless prohibited by local policy. If it currently has active clients, it SHOULD send a Reconfigure to each of them with the OPTION_NO_IPV4 included in the Option Request Option.

Example

- Dual-stack home router
  - Single WAN link
    - DHCPv4
    - SLAAC + DHCPv6
  - Single LAN link with multiple hosts
- On boot:
  - Assign 192.168.1.1/24 to LAN
  - Starts DHCPv4 server on LAN
    - Hands out 192.168.1.100-199 to clients
  - Starts IPv6 RA daemon and stateless DHCPv6 server on LAN

Example (cont.)

- Starts two provisioning processes in parallel: one for IPv4, one for IPv6
- In IPv6 process:
  - Router puts OPTION_NO_IPV4 in ORO in Request
  - Receives Reply with No-IPv4 level 2
  - Router aborts IPv4 process (if still running)
  - Deactivates all IPv4 functionality
  - Configures stateless DHCPv6 server to send the No-IPv4 option to LAN clients that request it
- Optimization: delay IPv4 provisioning process (10 seconds?) to avoid any IPv4 set up

Open Issues

- A legacy IPv4-only device connected to a network running in mode 2 (no IPv4 at all) will presumably keep retrying forever, e.g. sending DHCPDISCOVER messages endlessly.
  - Do we want a way to signal to that host that IPv4 will never be available?
  - But since that device was not updated for IPv6, it is doubtful that it would be updated to understand this new signaling. Could we reuse/overload some existing signaling that would have the same effect?

The last slide

- Questions? Comments?
- Next steps?