

Discourage IPv4 “proxyARP for everything” on IPv6-only networks

draft-yourtchenko-ipv6-disable-ipv4-proxyarp

Andrew Yourtchenko

Problem Statement

```
ayourtch-mac:~ ayourtch$ sudo tcpdump -ln -i en0 not ip6
Password:
tcpdump: verbose output suppressed, use -v or -vv for full protocol decode
listening on en0, link-type EN10MB (Ethernet), capture size 65535 bytes
14:56:22.789254 IP 0.0.0.0.68 > 255.255.255.255.67: BOOTP/DHCP, Request from 28:cf:e9:1d:22:f7, length 300
14:56:26.261680 ARP, Request who-has 91.189.89.134 tell 169.254.40.204, length 28
14:56:27.454032 ARP, Request who-has 91.189.89.134 tell 169.254.40.204, length 28
14:56:28.458465 ARP, Request who-has 91.189.89.134 tell 169.254.40.204, length 28
14:56:29.463602 ARP, Request who-has 91.189.89.134 tell 169.254.40.204, length 28
14:56:30.468971 ARP, Request who-has 91.189.89.134 tell 169.254.40.204, length 28
14:56:31.270648 ARP, Request who-has 199.59.148.82 tell 169.254.40.204, length 28
14:56:31.555933 IP 0.0.0.0.68 > 255.255.255.255.67: BOOTP/DHCP, Request from 28:cf:e9:1d:22:f7, length 300
14:56:38.185276 ARP, Request who-has 46.4.166.234 tell 169.254.40.204, length 28
14:56:38.387496 ARP, Request who-has 199.59.150.39 tell 169.254.40.204, length 28
14:56:39.394502 ARP, Request who-has 199.59.150.39 tell 169.254.40.204, length 28
14:56:40.398158 ARP, Request who-has 199.59.150.39 tell 169.254.40.204, length 28
14:56:40.433000 IP 0.0.0.0.68 > 255.255.255.255.67: BOOTP/DHCP, Request from 28:cf:e9:1d:22:f7, length 300
14:56:41.500506 ARP, Request who-has 199.59.150.39 tell 169.254.40.204, length 28
14:56:42.601809 ARP, Request who-has 199.59.150.39 tell 169.254.40.204, length 28
^C
15 packets captured
40 packets received by filter
0 packets dropped by kernel
ayourtch-mac:~ ayourtch$
```

Non-trivial portion of the
idle-state traffic

Strange ARP requests

Problem Statement

```
ayourtch — bash — 75x20
ayourtch-mac:~ ayourtch$
ayourtch-mac:~ ayourtch$ host stdio.be
stdio.be has address 188.40.136.148
stdio.be has IPv6 address 2a01:4f8:101:3245::cafe
stdio.be mail is handled by 10 mail.stdio.be.
ayourtch-mac:~ ayourtch$ telnet stdio.be 80
Trying 188.40.136.148...
^C
```

Why ?

- RFC3927, section 2.6.2

In the case of a device with a single interface and only a Link-Local IPv4 address, this requirement can be paraphrased as **"ARP for everything"**

Why ?

- RFC3927, section 2.6.2

In the case of a device with a single interface and **only a Link-Local IPv4 address**, this requirement can be paraphrased as **"ARP for everything"**

Only a link-local IPv4 address

- There's no other IPv4 address besides link-local

Only a link-local IPv4 address

- There's no other IPv4 address besides link-local
 - AND
- There is no IPv6 addresses that could have been used as an alternative

Suggestion: Update to RFC3927

"If the host has any interface with a global unicast IPv6 address assigned and any IPv6 route to any non-connected network (including default), then the host MUST immediately return an error rather than transmit any packet with a link-local IPv4 source address unless the destination is also an IPv4 link-local address."

Yes, but...

- You can turn off IPv4 manually
 - You can. I can. Not the users at large.
- You can turn off IPv4 via DHCPv6
 - Yes, but needs 2 vendors to implement.
- AF Sorting is the problem.
 - This is seen in more than one OS.
- It introduces the dependency between v4&v6
 - “ships in the night” is a leaky abstraction. Titanic.

Conclusion

- This is a problem that can be fixed in parallel with other efforts
- Tangible immediate results per host/OS
- Adopt as a WG item ?