Some problems observed in IPv6-only deployment

draft-song-sunset4-ipv6only-dns-00

Linjian Song, BII
Di Ma, ZDNS

Context of IPv6-only

“In order to fully transition the Internet to IPv6, individual applications, hosts, and networks that have enabled IPv6 must also be able to operate fully in the absence of IPv4. ”

----Charter for sunset4 WG

Real Use cases

- CERNET2 with IPv6-only Core
- T-Mobile US with IPv6-only access network
- IPv6-Only Internal Network in Facebook

Motivation

- Deployment of IPv6-only networks are impacted given IPv4-only or dual-stack transition scenarios
- To revisit and identify the implicit inertia of DNS which may hinder the IPv6-only deployment
- Hopefully propose a mitigation technique

DNS Proxy in IPv6 Only Network

- Tunnels make independent deployment for IPv6-only network
- DNS proxy approach is recommended in IPv6 only access network
  - Scenarios like DS-lite, lw4over6 and 464Xlat
  - The ISPs only provision IPv6 address to the DNS element via DHCPv6
  - Proxy help to connect to IPv4 world (NS servers)
  - IPv6 resolver and IPv4 resolver maybe configured in one device

Pitfalls of Proxy

- Proxy is not a full-function resolver
- Not implemented as the protocol described
  - Truncate all responses at 512 octets(or WAN MTU)
  - Without correctly setting the TC bit or remove it
  - Causing More retry over TCP
  - Either reject or black-hole any packet containing an OPT RR
  - Do not support extensions and hop-by-hop mechanisms
  - …

"To ensure full DNS protocol interoperability it is preferred that client stub resolvers should communicate directly with full-feature, upstream recursive resolvers wherever possible." --- **RFC6525**

Another cases of DNS inertia

- **Additional section in IPv4/IPv6 Environments**
  - Given the hard limit (512 bytes) in NS lookup response (Priming Exchange)
  - IPv4 glue and IPv6 glue of same zone are actually competing for the room of DNS additional section
  - Not all of the glue information can be included (RRset)

```
;; ADDITIONAL SECTION:
a.root-servers.net.   518400   IN   A      198.41.0.4
b.root-servers.net.   518400   IN   A      192.228.79.201
c.root-servers.net.   518400   IN   A      192.33.4.12
d.root-servers.net.   518400   IN   A      199.7.91.13
e.root-servers.net.   518400   IN   A      192.203.230.10
f.root-servers.net.   518400   IN   A      192.5.5.241
g.root-servers.net.   518400   IN   A      192.112.36.4
h.root-servers.net.   518400   IN   A      128.63.2.53
i.root-servers.net.   518400   IN   A      192.36.148.17
j.root-servers.net.   518400   IN   A      192.58.128.30
k.root-servers.net.   518400   IN   A      193.0.14.129
l.root-servers.net.   518400   IN   A      199.7.83.42
m.root-servers.net.   518400   IN   A      202.12.27.33
a.root-servers.net.   518400   IN   AAAA   2001:503:ba3e::2:30
b.root-servers.net.   518400   IN   AAAA   2001:500:84::b
```

DNS64 in IPv6-only Network

- DNS64 maps A RR to AAAA RR which works in T-Mobile US
- Some APP(Whatsapp) packs a large RR set that is probably close to 512 bytes on IPv4 full of A records to facilitate load sharing
- Android does not support EDNS0 [RFC6891]
- The DNS64 expands the A records to AAAA records, it exceeds 512 and results in a fallback to TCP

Thanks to Cameron Byrne who provide this scenarios in T-Mobile US

Beijing internet institute 工程中心

Comments Welcome!