DNS and RFC 8085 UDP Usage Guidelines

--- Avoid fragmentation, Again ---

draft-fujiwara-dnsop-avoid-fragmentation-01

Kazunori Fujiwara @ OARC 31

RFC 8085: UDP Usage Guidelines = BCP 145 (March 2017)

- BCP 145 specifies UDP usage guidelines including congestion control, message sizes, reliability, checksums, middlebox traversal, ECN, DSCP, ports
- Section 3.2. Message Size Guidelines
  - an application SHOULD NOT send UDP datagrams that result in IP packets that exceed the Maximum Transmission Unit (MTU) along the path to the destination.
  - An application SHOULD either use the path MTU information provided by the IP layer or implement Path MTU Discovery (PMTUD) itself [RFC1191] [RFC1981] [RFC4821] to determine whether the path to a destination will support its desired message size without fragmentation.

After then,

- Without cache poisoning attacks using IP fragmentation, RFC 8085 recommended to avoid fragmentation in DNS

- RFC 4035 and RFC 3226 need to be updated to avoid IP fragmentation

draft-fujiwara-dnsop-avoid-fragmentation-01 proposes

- UDP requestors and responders SHOULD send DNS responses with IP_DONTFRAG / IPV6_DONTFRAG
- The estimated maximum DNS/UDP payload size SHOULD be the actual or the default maximum DNS/UDP payload size
  - 1220 ≦ default maximum DNS/UDP size ≦ 1400
    - May be 1232
- Responders SHOULD compose UDP responses that result in IP packets that do not exceed the path MTU to the requestor
- Zone operator SHOULD consider small response size configurations
- How to retrieve path MTU value to a destination
  - getsockopt IP_MTU, IPV6_MTU on Linux
- Please review the draft