draft-ietf-dnsop-avoid-fragmentation

K. Fujiwara, P. Vixie

dnsop WG at IETF 110

Updates from IETF 109

- Submitted draft-ietf-dnsop-avoid-fragmentation-03, Nov. 23, 2020
  - includes proposed changes at IETF 109
- Submitted draft-ietf-dnsop-avoid-fragmentation-04, Feb. 22, 2021
  - Added reference to Geoff Huston's presentation "Measuring DNS Flag Day 2020" at OARC 34
  - Changed to use "default maximum DNS/UDP payload size" instead of "default path MTU value".
  - Added text about how to calculate maximum DNS/UDP payload size for each destination from the path MTU discovery result
  - Added text about "the server's path MTU to the Internet", and how to calculate the DNS server's default maximum DNS/UDP payload size

Recommendations changes at 03 04

3.1 Recommendations for UDP responders

- UDP responders SHOULD send DNS responses with IP_DONTFRAG / IPV6_DONTFRAG [RFC3542] options.
- If the UDP responder detects immediate error that the UDP packet cannot be sent beyond the path MTU size (EMSGSIZE), the UDP responder MAY recreate response packets fit in path MTU size, or TC bit set.
- UDP responders MAY probe to discover the real MTU value per destination.
- UDP responders SHOULD compose UDP responses that result in IP packets that do not exceed the path MTU to the requestor. If the path MTU discovery failed or is impossible, UDP responders SHOULD compose UDP responses that result in IP packets that do not exceed the default maximum DNS/UDP payload size described in Section 3.3.

3.2 Recommendations for UDP requestors

- UDP requestors SHOULD send DNS requests with IP_DONTFRAG / IPV6_DONTFRAG [RFC3542] options.
- UDP requestors MAY probe to discover the real MTU value per destination. Then, calculate their maximum DNS/UDP payload size as the reported path MTU minus IPv4/IPv6 header size (20 or 40) minus UDP header size (8). If the path MTU discovery failed or is impossible, use the default maximum DNS/UDP payload size described in Section 3.3.
- UDP requestors SHOULD use the requestor's (maximum DNS/UDP) payload size as the calculated or the default maximum DNS/UDP payload size.
- UDP requestors MAY drop fragmented DNS/UDP responses without IP reassembly to avoid cache poisoning attacks.
- DNS responses may be dropped by IP fragmentation. Upon a timeout, UDP requestors may retry using TCP or UDP, per local policy.

3.3 default maximum DNS/UDP payload size

- Default maximum DNS/UDP payload size for IPv6 is XXXX.
  - (Choose 1232, 1400, 1472 or other good values before/at WGLC)
- Default maximum DNS/UDP payload size for IPv4 is XXXX.
  - (Choose 1232, 1400, 1452 or other good values before/at WGLC)
- Operators of DNS servers SHOULD measure their path MTU to well-known locations on the Internet, such as [a-m].root-servers.net or [a-m].gtld-servers.net at setting up the servers.
- The smallest value of path MTU is the server's path MTU to the Internet.
- The server's maximum DNS/UDP payload size for IPv4 is the reported path MTU minus IPv4 header size (20) minus UDP header size (8).
- The server's maximum DNS/UDP payload size for IPv6 is the reported path MTU minus IPv6 header size (40) minus UDP header size (8).

Please choose the value and review

- Choose good "default maximum DNS/UDP payload size".
  - Default maximum DNS/UDP payload size for IPv6 is XXXX.
  - Default maximum DNS/UDP payload size for IPv4 is XXXX.

Source	IPv4	IPv6
RFC 4035 (MUST/SHOULD)	1220/4000	1220/4000
DNS Flag Day 2020	1232 ←	1232 1280-40-8
Access line case (with tunnel: PPPoE, v4 over v6, …) (Section 3.3 of this draft)	1400 ←	1400 1500-40-8-tunnel headers
Geoff Huston's "Measuring DNS Flag Day 2020" the prevailing MTU is 1500 between res. and auth.	1472 1500-20-8	1452 1500-40-8