DNS Transport over TCP - Implementation Requirements

J. Dickinson, S. Dickinson,
Sinodun Internet Technologies
R. Bellis, ISC
A. Mankin and D. Wessels, Verisign Labs

DNS Transport over TCP

- This is a -bis of RFC5966
- Now at -02 revision after much feedback
- In support of
  - Privacy efforts
  - Preventing amplification attacks
  - Packet size limitations

-01 vs -02 Major changes

- Re-structured Connection Handling section
  - Added Current Practice section
  - Made Recommendations section more granular
- Idle time discussion
  - Added definitions of Persistent connection and Idle session.
  - New text on recommendations for client idle behaviour.
  - Added statement that servers MAY use 0 idle timeout.
- Added more discussion on DoS mitigation in Security Consideration section.
- Re-stated position of TCP as an alternative to UDP in Discussion.

-01 vs -02 other changes

- Added more text to Introduction as background to TCP use.
- Move TCP message field length discussion to separate section.
- Updated text on server limits on concurrent connections from a particular IP address (soon to be network prefix).
- Added text that client retry logic is outside the scope of this document.
- Clarified that servers should answer all pipelined queries even if sent very close together.
- [Apologies] Glaring typo in first paragraph of Introduction

Historic TCP use

- Historically used only as a fallback option (TC=1) or for zone transfer.
- CLIENTS: Lack of clarity in earlier RFC’s, particularly wrt client behaviour.
    - Common for clients to do ‘one-shot’ TCP (inefficient).
- SERVERS: Server implementations were ‘basic’ in TCP connection management implementation (not much DoS mitigation).
- No DNS RFC discusses the term ‘persistent connection’

Persistence in 5966bis

- Introduced specific discussion of persistent connections
- Recognised and discussed in more detail the limitations of existing (compliant) implementations to manage persistent connections
- RECOMMENDATION:

  "To mitigate the risk of unintentional server overload, DNS client MUST take care to minimize the idle time of DNS-over-TCP sessions made to any individual server. DNS clients SHOULD close the TCP connection of an idle session, unless an idle timeout has been established using some other signalling mechanism."

Server TCP management

- Provide more detailed and specific recommendations to server implementors on how to mitigate TCP DoS attacks, along the lines of advice in e.g. HTTP drafts

- Hard to mandate behaviour in specs. Pointed to best practice and general guidance. (HTTP has a wealth of experience handling persistent TCP.)

Pipelining / OOOP

- Concern on mailing list about clients unable to handle OOOP
- However consensus seemed to be that all existing clients that performed pipelining also handled out-of-order responses so and additional signalling mechanism was an unnecessary overhead.
- Agreement that most modern server implementations support pipelining

Re-try on failure

- Out of scope for this document. Left to implementors.

5966bis

- Thank working group for discussions
- Hope addressed majority of comments
- Would like to progress to Last Call in near future