ANY

Joe Abley
Ólafur Guðmundsson
Marek Majkowski

https://datatracker.ietf.org/doc/draft-jabley-dnsop-refuse-any/

Why is ANY query bad?

- Amplification

- Information Leak

- Unreliable

- Expensive

ANY uses

1. Debugging: i.e. Human asks and parses answer

2. Probabilistic Optimization: trying to get one or more answers in one query, Program asks then falls back on direct queries if answer is not useable

3. Misunderstanding by a programmers

4. Amplification attacks

Draft Goal:
be simple and non-disruptive

- Assumption #1 ANY != ALL

- Assumption #2 RCODE != 0 is bad

- Assumption #3 Resolver should cache

  - Something is better than nothing

- Assumption #4 Human should Understand

- Document how to avoid returning big answers

Resolver behavior

- Popular Resolvers interpret RFC2181 to mean

  - Answers returned in ANY query can not be used to answer for that type!!!

    - This makes probabilistic optimization useless on cold cache

Approach

- Return one RRSET

- DNSSEC signed if needed

- Each implementation can chose what they want

- fake answer is fine as long it is harmless

  - HINFO suggested as it can be displayed by deployed code

Next steps

- Draft has been adopted

- Almost Ready for WGLC ?

- Number of implementations deployed