Opportunistic recursive to authoritative: a protocol proposal

draft-pp-recursive-authoritative-opportunistic

Paul Hoffman

DPRIVE working group

Interim meeting, 2021-01-26

The proposal

- Use case
- How resolvers can enable this
- How authoritative servers can enable this
- Future possibilities for the draft

Use case

- Recursive resolver operators who are happy to use TLS encryption with authoritative servers if doing so doesn’t significantly slow down getting answers
- Authoritative server operators that are happy to use encryption with recursive resolvers if it doesn’t cost much
- Don’t fail to serve queries that would have worked over classic DNS on port 53

There will be extra costs when deployed

- It’s OK that there is an additional cost for this
  - Extra round trips to establish TCP for every session
  - Extra round trips for TLS establishment
  - Greater CPU use for TLS establishment
  - Greater CPU use for encryption after TLS establishment
  - Greater memory use for holding TLS state

How resolvers can enable this

- Use a cache that tells what is known about each authoritative server’s transport capabilities
  - Only do DoT if the cache says so
  - Fill the cache out-of-band
- Authenticate only if it is useful; otherwise, don’t authenticate or ignore the result if you have to authenticate
- So far, there is no agreed-on reason to authenticate in this protocol, so maybe we can just delete it

How authoritative servers can enable this

- Turn on TLS˜
- Maybe use a certificate that might be useful for clients that authenticate, or maybe just use a self-issued certificate
- Serve normally

Future possibilities for the draft

- WG adoption?
- Add TLSA records for another route to faster discovery for the cache
- On optional authentication, either:
  - Define where authentication during opportunistic recursive-to-authoritative is useful, and write more about how to handle authentication
  - Delete everything about authentication and leave it to a possible proposal for always-authenticated proposal