Encrypted DNS server redirection (EDSR)

Tommy Jensen, Microsoft
John Todd, Quad9

IETF 115

Problem

Anycast DNS provides the convenience of a single set of IP addresses or hostnames that work for everyone, but at a cost:

- Extensive BGP knowledge (staff) and OpEx to build/maintain
- Anycast configs more difficult or impossible in less-dense edge network environments
- Anycast cost tends to be a gating factor that limits less-resourced operators
  - Result: **Anycast models tend to reinforce centralization momentum**

Problem

Without anycast, client routing edge cases poses issues:

- Clients which erroneously end up talking to a non-geolocated server (distance)
- Clients which erroneously end up talking to a server in the wrong policy zone
- Unicast servers that need to shed/distribute excess traffic load
- Anycast has no good way of differentiating service profiles – everyone in the same pot

Requirements

- MUST NOT reduce security from the original connection when redirecting
- MUST NOT break compatibility (redirecting to server the client cannot connect to)
- SHOULD support encrypted DNS generally, not a specific subset
- SHOULD NOT introduce any more perf cost than absolutely necessary

Proposal

Reuse the DDR mechanism – use designations as redirections

- When connecting to an encrypted DNS server, start with resolver.arpa query
- If designations are returned, treat then as redirections
- New server identify verified by name, not by IP address
  - Unlike DDR, original query is encrypted and content is trusted
- Redirection valid for lifetime of SVCB TTL

Proposal

Example: client is configured to use doh-sydney.site.example as a DoT server

- Client sends SVCB query for resolver.arpa to doh-sydney.site.example
- Server returns doh-paris.site.example SVCB and additional A/AAA records
  - Because it sees the client is based in France, not Australia
- Client makes a new connection to doh-paris.site.example
- The TLS connection is validated using the “doh-paris.site.example” name
- If successful, doh-Sydney connection is closed

Proposal

Considerations

- Server MUST NOT redirect clients to servers which do not (at least) support the encrypted DNS protocol and IP address family it sees the client using
  - This ensures clients do not get redirected to a server they cannot communicate with
- Deployments should be mindful of avoiding long redirect chains

Alternatives considered

- HTTP 3xx
  - Not generic across protocols
  - Introduces per-query overhead (where EDSR introduces per-connection overhead)
- Alt-SvcB
  - This limits the responsiveness of redirections (as a property of the server's domain name rather than a specific connection)
  - Redirection as a concept significantly different than an alternative service

Conclusion

EDSR enables encrypted DNS server redirection by reusing DDR mechanics, which…

- Provides a one-size-fits-all solution
  - Works for any TLS-based encrypted DNS protocol, including DoH, DoT, and DoQ
- Encourages decentralization by leveling the playing field
  - Eliminates the need to support anycast infra to avoid complex, localized configuration when deploying globally – "first" server becomes a rendezvous
- Reuses existing records and mechanics

Questions?

Seeking WG adoption