DNS OVER HTTPS (DOH)

Performance Implications & Risks

2018 Workshop on Internet Economics (WIE)

December 2018

COMCAST

DOH – THE PROTOCOL

DEVELOPED TO COUNTER PERCEIVED THREATS

- Hostile governments & political environments:
  - Surveillance by listening on the wire
  - DNS modification to prevent access to politically-sensitive content
- Malware and attacker-based DNS response modification:
  - Attacker on the wire between stub and resolver
- Commercial use of DNS data:
  - Passive collection, sharing, and monetization of DNS queries by DNS operators (e.g. ISPs, WiFi hotspot operators)
  - Active DNS modification for NXDOMAIN redirection

HOW IT WORKS

- Encrypts over-the-wire communications between the stub (client) and recursive server, via TLS
- Uses TCP/HTTPS as a protocol rather than UDP/TCP port 53 DNS
- Brings in much richer HTTPS-based client capabilities (including fingerprinting/tracking) compared to simpler and more compact UDP/53 DNS protocol

SO FAR SO GOOD!

COMCAST

DOH – THE IMPLEMENTATION

THE STUB IS THE BROWSER OR MOBILE OS

- Google Chrome (61%), Apple Safari (15%), Mozilla Firefox (5%)
- Android (75%), iOS (22%)

"TRUSTED RECURSIVE RESOLVERS"

- Each browser / OS appears likely to choose its own default resolver
- The client software will turn it on by default (so the transition could happen rapidly is 2 or 3 actors implement)
  - Google: Google Public DNS (assumed)
  - Mozilla: Cloudflare (announced)
  - Apple: Unknown

IF JUST GOOGLE AND MOZILLA MOVE

- 75% of the world's mobile devices switch from current DNS to centralized DoH
- 2/3^{rds} of the world's web browsers switch from current DNS to centralized DoH

UH OH…

COMCAST

RISKS

DRAMATIC CENTRALIZATION OF THE INTERNET'S MOST WIDELY DISTRIBUTED PROTOCOL

- A centralized commercial authority decides unilaterally what performance/security tradeoffs to make for users, as opposed to use—driven tools like VPNs
- **75%** of mobile-based DNS traffic to one US-based commercial provider & **2/3** of browser-based DNS traffic to two US-based commercial providers (61% to just one)
- Tantalizing attack target: hit 2 or 3 operators and take down the global DNS via BGP hijack, DDoS, compromise of internal tools/systems, compromise of an admin account from 1 – 2 dozen sys admins
- Alluring surveillance target: hit 2 or 3 operators to target for surveillance / collection
  - Via legal (incl. US NSL) or extra-legal means, including on the wire, in the data center, in the hardware
- Data monetization bonanza: just 2 operators have full history data
  - 2 operators will have full history data on 2/3 of the global Internet where they have little/none today
  - Reidentification likely to be trivial & trackability moves to the device/individual level, just like other HTTP tracking

NO MORE LOCAL POLICY EXPRESSION IN EACH NETWORK CONNECTING TO THE INTERNET

- DNS-based parental controls & malware detection (ISP, campus/EDU, enterprise, government), RPZ
- Split DNS (i.e. internal-only names), including private corporate names; ISP provisioning and Hot-spot splash screens
- Passive DNS security tools

COMCAST

MEASUREMENTS TO DATE

SEVERELY LACKING

- For such a potential large-scale change, significant and dependable measurements are required, including peer review and community consideration of the results and implications.
- Measurements to date come from only from Mozilla:
  - 25,000 self-selected users of the Firefox “nightly build” (users that opt-in to test new features)
  - These self-selected users may or may not measure enterprise, campus, and ISP breakage risks
  - No idea if local factors such as WiFi, traffic concurrency, or other issues confounded the measurements
  - “Most” queries (whatever most means) were said to be 6 ms slower. But this only measures DNS query response time, not the time to fetch the destination content and whether that was fully localized via a CDN – in essence it was slower to get AN answer and unclear if it was the BEST / MOST LOCAL answer.
  - This placed no significant load on the end resolver, so was not a representative load test. This is a concern as the resolver load on a per-query basis is likely to be much higher for DoH vs. UDP/53 DNS. For comparison, the Comcast DNS resolvers receive over 500 billion queries per day. The infrastructure to handle 25,000 users for a few hours is in no way comparable to billions of queries, so no server-side scaling conclusions can be drawn.

COMCAST

MEASUREMENTS NEEDED

BEST ANSWER VS. ANY ANSWER, OPEN DATA, BETTER CONTROL OVER VARIABLES

- Control the end point so as to avoid the influence of confounding local factors such as WiFi, traffic concurrency – using something like RIPE Atlas or (FCC MBA) SamKnows probes
- Distinctly compare a control (ISP DHCP-issued DNS resolvers) vs. 3rd party public DNS resolvers vs. DoH resolvers
- Test query response time (QRTT) for each of these resolvers
- Test HTTP content retrieval time for each of these resolvers – ensuring that the queried names include names that are CDN-based and likely to be most localized (most popular content – not long tail)
  - This has never really been done at scale, even for ISP resolvers vs. 3rd party public resolvers
- Compare all results by network/DNS resolver/DoH resolver and by geography (i.e. state, country, continent, time zone)
  - Possible some networks and resolvers will be better than others, results may vary geographically
- Compare content retrieved (esp. look for failures such as broken geo-fencing, etc.)
- Make resulting data publicly available for study
- There are many opinions about performance – but little data exist – data can help

COMCAST