April 2019
Geoff Huston

More DOH

It seems that the previous article on DOH has generated some reaction, and also there is some further development that should be reported, all of which I'll cover here.

> The previous article on the reactions to DOH at IETF 104 can be found at https://www.potaroo.net/ispcol/2019-04/angst.html

Default DOH

First, with respect to DOH as a default setting in browsers I had noted: "If a browser chooses to use DoH [as a default setting] then there is little that the platform or the network can do to prevent it. If a browser has installed DoH support, then control over the DNS name resolution function has passed from the user to the browser provider, and rather than being an esoteric function enabled by a handful of users, it becomes a "mainstream" service used by potentially billions of end users."

On April 9 Mozilla announced its plan to enable DOH by default in the Firefox browser [https://mzl.la/2P8ypMk], committing to an earlier informal description of Mozilla's plans that were outlined by Mozilla's Eric Rescorla at the end of March [http://bit.ly/2UhltVl].

To place this announcement into a broader perspective it should be noted that the market share of Mozilla's Netscape browser, while large, is by no means dominant. The *statscounter* site reports a market share of 4.69% for Firefox in March 2019 [http://gs.statcounter.com/browser-market-share]. This implies that these moves by Mozilla are not intrinsically all that significant in terms of the profile of the larger Internet and the average Internet user. A major concern with this announcement is that the move by the Firefox browser to make DOH the default means of DNS name resolution is a precursor for similar changes to the Chrome browser. Chrome is definitely the dominant browser in today's Internet ecosystem, with some 62.63% market share according to *statscounter*. If Chrome were to use a default setting that pushed all its DNS name resolution activities to a Chrome-selected DOH server then the implications for the DNS are very significant.

In APNIC's in-progress study of the use of open DNS resolvers (such as the services operated by Google, Cloudflare and Quad9, among others) the total market share of these open DNS resolvers encompasses approximately 20% of the Internet's user population, while most other users use ISP-provided DNS resolvers. If some two thirds of the Internet's user base had their DNS queries redirected to one, or even just a handful of these DOH-based DNS services, as would likely be the case were Chrome to enable DOH by default, then this would be a major change to the Internet. The open source DNS resolver effort, populated at present by ISC's Bind DNS resolver, NLNet's unbound, CZNIC's KNOT and PowerDNS, among others, would probably not survive this scale of change, and the risk is that the open DNS itself may not survive. The privacy implications are also serious. The DNS would still create a rich vein of information about each user's activities, but the parties who are privy to this information flow would change, creating potential new winners and losers in the marketplace of surveillance capitalism.

Will the other browsers follow Mozilla's lead with DOH enabled by default? The experience so far would support a "yes" answer. Browser vendors have been enthusiastic to integrate changes to their platform that decrease page load times and equally keen to integrate changes that protect the browser's activity against various forms of surveillance. DOH does not necessarily make DNS resolution quicker, although it does put the browser in more control over its use of the DNS and allows the browser to control its own local DNS cache. But, of course, DOH plugs a critical DNS information leak in the current browser architecture. Third party observers can infer browser activity by looking at the browser's DNS query stream. DOH prevents any such observation in either the user's platform or the local network. So "yes" is a likely answer to this question.

Bad DOH

By pushing client-side DNS queries into HTTPS the Internet itself has effectively lost control of the client end of DNS, and each and every application, including the vast array of malware, can use DOH and the DNS as a command and control channel in a way that is undetectable by the client or client's network operator.

Much of today's malware containment frameworks, including DNS firewalling, are rendered useless by DOH. Whether or not the browser has DOH enabled by default, applications can generate DOH requests for DNS resolution in a manner that bypasses today's DNS-based malware containment mechanisms. As has been recently observed on a DOH-related mailing list: "Pandora's box is now open and DOH has escaped, and there seems to be little we can do about it now. The times they are a changing."

Surveillance DOH

I raised the question in the previous article: "Have we now provided the private surveillance framework with a whole new trove of personal data to mine by ruthlessly exploiting the DNS in a manner that is entirely out of sight? Once the browsers and even the apps direct their name queries through encrypted channels to resolvers operated by the same browser and app providers, then have we dealt a body blow to any efforts to safeguard personal privacy on the Internet?"

DOH service operators have a clear view of the end user. The end-to-end encryption mechanism of TLS implies that the query is being passed from the end application, rather than via any DNS intermediaries. This allows a DOH service operator to assemble individual user activity profiles, and the concern here is that such individual profiles have considerable value in the online advertising market. Open DNS resolver operators offer a "free" service, as users are not charged fees to resolve DNS names. Would the temptation to fund this free service by monetizing such individual profiles prove to be overwhelming?

I should clarify this speculation further, as it has generated comment that the major open DNS resolver services, who appear to also be the current DOH service operators, already operate with the terms of clearly stated privacy policies, and the ruthless exploitation of personal data appears to be entirely out of scope of such policies, obviously. What do today's open DNS resolver operators that provide a DOH service say in their privacy policies?

The Quad9 Open DNS resolver has a privacy policy which states: "We share anonymized data on specific domains (such as domain, timestamp, geolocation, number of hits, first seen, last seen) with our threat intelligence partners. Please note that this information does not contain source IP information or any other identifier that would directly identify the end user or their organization." [https://www.quad9.net/privacy/]

Google state that: "We don't correlate or combine information from our temporary or permanent logs with any personal information that you have provided Google for other services." [https://developers.google.com/speed/public-dns/privacy].

Cloudflare are also quite explicit in this area: "Cloudflare's business has never been built around tracking users or selling advertising. We don't see personal data as an asset; we see it as a toxic asset. While we need some logging to prevent abuse and debug issues, we couldn't imagine any situation where we'd need that information longer than 24 hours. And we wanted to put our money where our mouth was, so we committed to retaining KPMG, the well-respected auditing firm, to audit our practices annually and publish a public report confirming we're doing what we said we would." [https://blog.cloudflare.com/announcing-1111/]. Engaging a third party to audit a private enterprise's claims of privacy practices seems like a good move, as without some form of visible accountability, the company's privacy policies are, ultimately, just words.

These open DNS providers appear to have a clear view of user concerns over personal privacy. Their privacy policies implicitly acknowledge that the DNS query stream could be used to provide insights into the personal profile of users and assert that that have no such intent to do so. Such noble intentions to operate a free public service and refrain from any form of monetization of the service are entirely laudable.

However, from an historical perspective these undertakings appear to be unrealistic and unsustainable. We should remember the events of a century ago with Theodore Vail and the Kingsbury Commitment in 1910 in the United States. His key commitment was a profession of noble intent to enrichen the public space. AT&T was to be an "enlightened monopoly" that served the public in close cooperation with the state while at the same time serving the interests of AT&T shareholders. His view of the telephone service as a privately-operated public utility, is "at once the most sympathetic and scariest element of his vision. Vail saw no harm in, and indeed believed in, giants, so long as they be friendly giants. He believed power should be beneficently concentrated, and that with great power came great responsibility." [Tim Wu, "The Master Switch: The Rise and Fall of Information Empires," Atlantic Books]

As we observe the aggregation of this critical part of the Internet's infrastructure in the centralization of the DNS, it cannot be ignored that these grand statements of respect for the public interest and undertakings that safeguard personal privacy sound scarily similar to the espoused public benefactor vision of AT&T in 1910 as it embarked on a course of establishing a national monopoly. But it is perhaps not today's operators and today's commitments that should concern us, but where this may lead. Again, quoting Tim Wu: "[Theodore Vail] presents us therefore with a challenging figure: an unabashed monopolist, but a benign one, who lived up to his own ideals of enlightened despotism. The fault in this arrangement therefore lay not so much with Theodore Vail as with the men who would succeed him." [Tim Wu].

Perhaps the same is true of these current undertakings relating to protection of personal privacy and their perception of the greater public interest. Over time these earnest undertakings in the provision of free services may well be eroded by the inevitable pressures that every private enterprise is prone to, namely those of paying the bills and maximizing shareholder value. Once the DNS is placed under an all-encompassing shroud of deep encryption then both good and dark deeds will be undetectable.

Name Space DOH

We have reached a very odd place with today's Internet.

The response to running out of IPv4 addresses has been the massive use of address sharing practices. We've crammed more than 20 billion devices into some 3 billion IPv4 addresses. Yes, if we ever get to the other end of this protracted transition to IPv6 there is a vague prospect that we will be able to restore address integrity, but this is somewhat unlikely. But right now, addresses are semantically confused. At best in this environment of intense address sharing, IP addresses are merely ephemeral session tokens. It appears that what holds the Internet together as a single network is a single coherent name space.

But will this still be the case when the name resolution function, the critical element of the name space, is shifted behind an opaque shroud? Will the name space maintain its coherency and consistency when

there is no ability to oversee the entire name space? We have already seen efforts to use the DNS to steer users to the closest content location by tailoring the response to suit the querier, but with DOH it is possible to go much further in customizing views of the name space based on the identity and location of the end user and the application that they are running. What becomes of a coherent name space when the resolution of a name depends on who is making the query?

These are indeed interesting times for the Internet.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net