SUBJECT: DOH comments
FROM: jbash aka John Bashinski <[redacted]
TO: [redacted]
DATE: 20/11/2020 01:13
ATTACHMENTS (20201120-011338-0000030): [redacted]

At the very least, DoH should never be on by default for anybody, ever. You can make it opt-in if you must. It's still a bad idea to spend time on it at all, but at least you won't be breaking people who didn't ask to be broken.

Your specific questions

"Respecting privacy and security"

> To what extent can our requirement be shortened further while
> allowing providers sufficient data to operate the service?

There is no legitimate, pressing need for a DNS resolver to keep any logs of any kind, period, not for 24 hours and not for 24 seconds. Your policy should be to completely forbid all retention of query data.

It's OK to keep *statistics* about total query volume and which high-level zones are seeing a lot of queries. It *might* be OK to keep statistics about which of your BGP peers are sending you the queries. It's not OK to keep anything more granular than that.

> Are there exemptions that should be allowed by the policy for
> additional data collection in emergency circumstances?

No. That's a horrifically bad idea, because

1. The resolver *will never have any way of authenticating whether an "emergency" actually exists*.

2. It invites building an infrastructure that allows for the creeping expansion of what constitutes an "emergency".

"Emergency exceptions" are basically a vehicle for handing power to bad guys.

"Online safety"

Your whole concept of "internet safety" is incoherent. You conflate...

+ "Illegal content". This basically means censorship against the user's will, and presumptively against the user's interests.

... with ...

+ "Harmful content". This at least *connotes* some intention to
serve the user's interests... but leaves open the vast question
of what those interests actually are, and who's supposed to
decide what they are.

> Our current policy states that the provider operating the resolver
> should not by default block or filter domains unless specifically
> required by law in the jurisdiction in which the resolver
> operates. How, if at all, should this requirement change to address
> legally required blocking in other jurisdictions?

It should not be changed. You should not be making any change
that makes it any easier for anybody to block any name at the
DNS level, ever.

> What harmful outcomes can arise from filtering/blocking through the DNS?

Whoever has the power to block has the power to prevent communication
between any parties they want, thus undermining the whole point of the
Internet and of your software. What worse harmful outcomes do you
want?

The failures will presumably be confusing ones where names don't
resolve, but the user isn't told why.

There are also going to be a bunch of random software bugs and weird
behaviors that get triggered when a name that "should" resolve fails
to do so, but I don't think any of them will end the world. The
biggest problem is with the *intended* functionality.

> What more rights-protective and technically effective means of
> protecting users from illegal and harmful content exist beyond
> DNS-based blocking?

The ones that are already in wide use in Firefox are *all* more
rights-protective and more technically effective. DNS blocking
brings *nothing* to the table.

It's trivial to have a local domain or URL blocklist, or even a remote
one, without wiring it into the DNS infrastructure of the
browser. It's also not hard to do content-based blocking or even
behavior-based blocking.

You already *have* at least one or two block lists built in. There are
also a couple of very widely used extensions that can already do
everything DNS blocking could do and more.

You have an add-on system. Let users opt into whatever filters they
want. You do not have to involve yourselves in that process, and
you should not be spending time on it.

In fact, the "safety" stuff you already have is mostly annoying,
privacy-damaging misfeatures that have to be turned off at every new

browser install. Don't make it any worse by screwing up the DNS as well.

> How could we ensure effective transparency and accountability in
> situations where TRRs engage in legally required blocking practices?
> (For example: publicly available transparency reports with blocked
> domain names by country.)

You can't, because the governments of those countries will outlaw distributing such lists, if they haven't already done so. This is another reason to minimize the number of governments you enable to block things.

> What governance, process, or audit requirements should be required
> of parties that maintain and create block lists? For example, what
> complaint and redress processes should exist?

Resolvers should not be doing any blocking whatsoever on their own initiative, and it's not reasonable to expect a resolver to process complaints about government orders to block things. Therefore you should not require a complaints process.

> How can we best present information about opt-in filtering endpoints
> to end users (e.g., for malware blocking or family-friendly
> blocking)?

There are a nearly infinite number of more important browser features to add before you worry about that. Just don't do it. Again, you have an add-on system. Put the energy into improving the APIs for extensions.

"Building a better ecosystem"

> How can deployment of DoH help to increase trust in Internet
> technologies in your region?

It can't and won't. It will just break things, confuse people, and spread paranoia.

In any region where it would actually do any good, even for the most naive user, it will be blocked by packet filtering, subverted by what will appear to your software as "enterprise policy", or obviated by DPI and SNI sniffing.

> What exploitations of the DNS in your region could DoH protect
> against?

Obviously it protects against local spoofing. It just does a worse job of end-to-end protection than, say, DNSSEC.

> What are the best ways to gain global adoption/support of the DoH
> standard amongst ISPs and DNS providers?

Mass lobotomies? Seriously, all these questions just *assume* that DoH

is a good thing that should be promoted. It's not. DoH is a bad idea that should die.

> Are there specific DNS use cases for which you think DoH would
> provide particular security and privacy value (e.g., when users
> connect over free public WiFi hotspots)?

None that could not or should not be done better in the OS and the rest of the network stack. If I can't trust the free hotspot's DNS for Firefox, then I can't trust it for any other application either.

> Although Firefox disables DoH when it detects that enterprise
> policies arein place, are there other situations in which deployment
> of DoH might cause technical or operational challenges (e.g., mobile
> networks, NAT64 and DNS64)?

Various "local name resolution" hacks will probably break, leading to things like not being able to manage your printer. Attempts to work around this by exempting ".local" or whatever will work imperfectly.

People who've already opted in to other DNS bad ideas like OpenDNS will probably see unexpected behavior.

I strongly suspect that Firefox with DoH on will ignore the elaborate local resolver configuration on my laptop, because it's not "entrprisey" enough and won't get detected by whatever heuristics you have. The same applies for any vaguely sophisticated user. The real point is that nobody even KNOWS what will break. People will be tearing their hair out because some DNS thing "works" in their other applications or their debugging tools, but not in Firefox, or vice versa.

... and a bunch of the usual suspects will surely block it intentionally using various weird firewall tricks. If they do that, and if you enable DoH by default, then you have no reasonable response. If you stubbornly keep demanding to use DoH, you will have no DNS resolution. If you give up and fall back to normal DNS, then DoH can be subverted by the very people you most want *not* to subvert it.

DoH generally

DoH was a bad idea from day one. It doesn't matter if your resolvers are "trusted". It doesn't matter if your resolvers are *incorruptible*. It's still a bad idea.

+ You are not increasing privacy or reducing leakage via DNS. You're circumventing system policy, and thereby making it harder to get DNS privacy globally.

The right place for decisions about DNS is in the OS, not in every random client application. If I don't want to use my ISP's resolvers (and I don't), *my system lets me do that*. I don't want you

overriding it.

Crude attempts to detect "enterprise" policy are not a substitute for *just using the resolver you are given*. That kind of guesswork is technically unreliable. It's probably easy to subvert; there's no real way to tell whose policies should be believed. And there's something *deeply wrong* with the whole idea that "enterprises", rather than users, should set policy to begin with.

You aren't going to deny ISPs much, anyway, because they have access to the actual data traffic. They can and will find out where people are going using not only IP address snooping, but also HTTP header snooping, SNI snooping, and other forms of DPI. All you're really doing is creating an *additional* point at which information can be stolen.

If you really care about DNS information leakage, start by turning off prefetch by default.

+ You are not really addressing integrity. You're creating a patchwork path from the zone authority to the user, with some parts protected and some not.

If you really care about integrity, use DNSSEC (and DANE). Or, better yet, let the OS resolver worry about DNSSEC for name resolution, and *only* involving yourselves for DANE purposes.

+ You are not preventing censorship. You're handing total censorship authority to whatever governments or other forces can put pressure on the very large, centralized coporate players who serve as "trusted resolvers". I see that you're now looking at how to give the veto to *more* governments, with "address legally required blocking in other jurisdictions".

If you really care about censorship resistance, look into parallel support for decentralized, permissionless, non-DNS name resolution systems, instead of spending your time providing *more* infrastructure to censor the DNS.

+ You are not promoting "safety".

Nobody wants you, or your resolvers, to be in charge of determining what's "safe". It's not a question of what policy you should have; the deeper issue is that *it's not your policy to set*. Maybe you're forced to take on management of the trusted CA list. Taking on policy for DNS is an *unforced* error. You have enough trouble delivering a working Web browser.

+ You are MOST DEFINITELY NOT "respecting the technical architecture of the Internet".

+ You're abandoning fate sharing, and damaging resiliency with unneccessary centralization.

+ You're ignoring separation of concerns.

+ You're adding the unnecessary complexity of HTTP into the comparatively simple DNS protocol. Seriously. HTTP???

+ You're creating a bunch of special cases that will create bugs and cause incomprehensible behavior.

... and you're siphoning off attention from architecturally superior alternatives, by which I mean basically every other proposal for DNS encryption or integrity. Many of those are bad, but DoH is the bottom of the heap.

The right answer is to kill the whole thing. I know you won't do that, but at least don't force me to remember to manually turn it off every time I install your browser, or to spend a bunch of time tearing my hair out every time I forget to do so.

Nobody wants this. Why do you insist on inflicting it on us?

-- jbash