Consultation Response

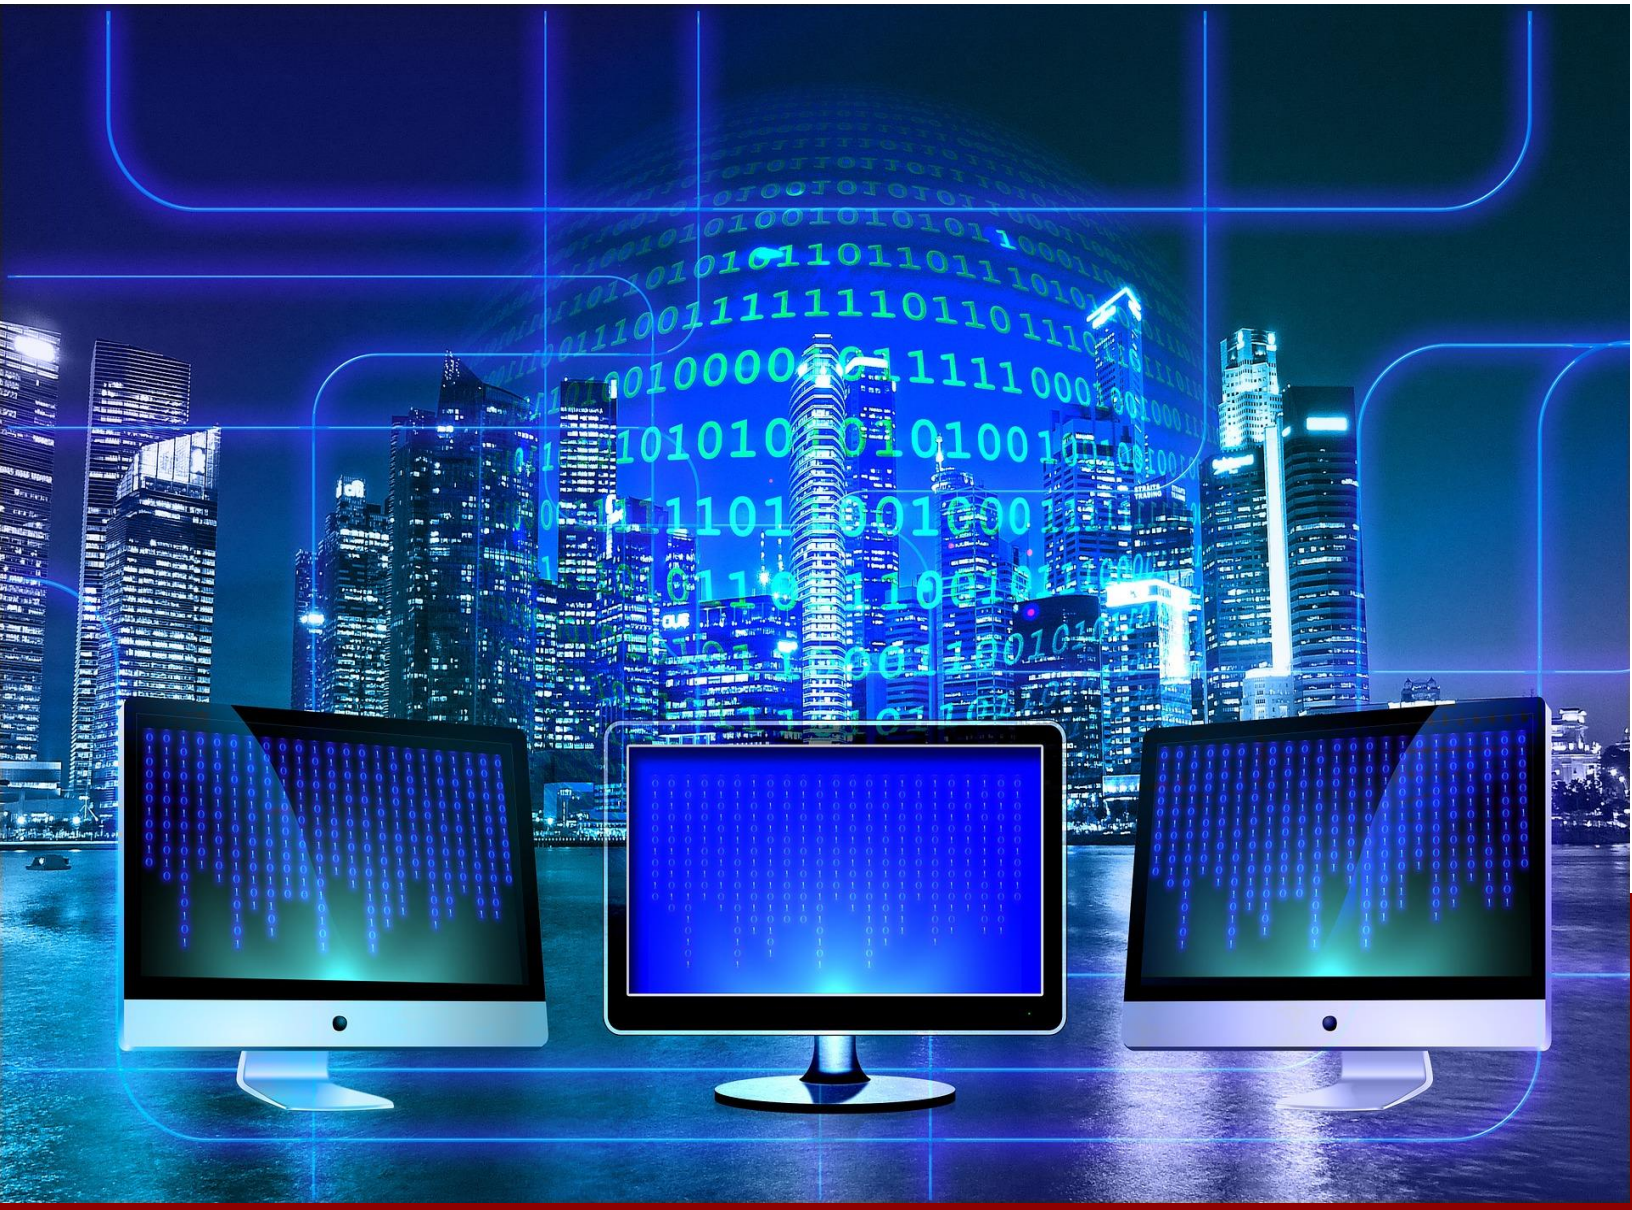

Mozilla DNS over HTTPS and Trusted Recursive Resolver

Andrew Campling

20th January 2021

419.Consulting

(This page is intentionally blank)

Cover image by Gerd Altmann from Pixabay

Introduction

The Following document is a response by 419 Consulting[1] to the public consultation by Mozilla Corporation regarding its Trusted Recursive Resolver Policy. As a general comment, the encryption of DNS to improve user privacy is a laudable aim, however, care must be taken to ensure that the method of implementation does not create new harms that outweigh the benefits being offered. The issues to be considered include:

- Whether the bypassing of network-based malicious content filtering leaves users exposed to harmful exploits;
- Whether the bypassing of illegal content blocking causes harm to the subjects of that material (in the case of child sexual abuse content), economic loss (in the case of copyright-infringing material) or other, negative consequences;
- Whether the bypassing of adult content filtering leaves children needlessly exposed to inappropriate material;
- Whether centralisation of DNS resolution reduces user choice and/or harms the resilience of the Internet;
- Whether centralisation of DNS resolution leaves users even more exposed to surveillance capitalism and/or increases their cybersecurity risk by creating larger pools of data.

The above list is not intended to be exhaustive but to illustrate that the choices being made are not as straightforward as may first appear.

A second general issue worthy of consideration is the difficulty in obtaining GDPR-level consent from a user to a change in the DNS resolver. Noting that the vast majority of end-users have no awareness of the existence of DNS, its function nor how it works, a simple dialogue box asking whether they wish to change their DNS settings is unlikely to be deemed sufficient to gain informed consent from the perspective of GDPR, and the other types of GDPR-level consent are likely to need additional information too. This needs to be addressed.

Additionally, any interruption in a cloud-based DNS service, as has happened to, for example, Cloudflare's offering in the last few months, is likely to cause confusion, resulting in service calls to the user's ISP rather than the resolver operator. Any curated offering like that provided through the Mozilla TRR should include provision for multi-channel customer service in the relevant local language(s) in the event of any issues arising.

Finally, the current TRR approach does not take into account the potential resolver discovery standard being investigated through the IETF's ADD working group. Nor is there any effort taken to determine whether the user's current DNS service is encrypted or, if not, if the resolver operator has an encrypted offering. Consideration should be given to including these to provide more choice and respect existing user settings.

To aid understanding, the text from the consultation document has been included in the following pages, with any response to the points raised inserted immediately after the relevant section of text in ***red italics.***

[1] Please direct any queries to Andrew.Campling@419.Consulting

Questions for Comment

Mozilla Comment Period on DNS-over-HTTPS Implementation

We are seeking comments in four areas. Firstly, we seek general feedback with respect to our TRR policy and its relation to different regions. We also seek to crowdsource helpful input in three specific areas related to product roll-out in new regions, which will help us maximise the security- and privacy-enhancing benefits of default-on DoH for more users.

Some of the comments below refer to the European Resolver Policy[2]. This has been developed by representatives from across the telecoms and tech sectors, with input by civil society, governments and regulators, to provide a robust set of GDPR-compliant policies that can be adopted by resolver operators and others that seek to offer their services in Europe. The policy will be launched in early 2021, with some of the responses to this consultation incorporating excerpts from the current draft text.

General comments regarding our TRR policies

DNS over HTTPS (DoH) brings the benefits of transport-level security to DNS queries and responses. Building on this foundation, Mozilla partners with selected DNS providers who join our Trusted Recursive Resolver (TRR) program to ensure even stronger privacy and security guarantees for Firefox users. This means that DoH look-ups in Firefox are routed to DNS providers who have made binding legal commitments to adopt extra protections for user data. Our TRR policy sets strict conditions regarding the handling of DNS data; in particular, it establishes limits on data collection, use, and retention, limits on filtering and blocking without user consent, and transparency regarding data handling.

Consistent with the transparent practices and commitment to openness that Mozilla is known for, we welcome general feedback on our TRR policy and its relevance for particular regions in different parts of the globe - what benefits it may bring in terms of privacy and security, and what local considerations we should be conscious of in different regional contexts.

A major weakness of the approach taken by the current TRR policy is that it leads to a significant centralisation of Internet infrastructure, specifically DNS resolution. In the case of the US market, only three resolver operators are currently approved which is a major, undesirable reduction in user choice. This centralisation has negative consequences, both in terms of infrastructure resilience and by creating an attractive target for malicious actors of all types, including those with state support.

[2] See www.EuropeanResolverPolicy.Com or email Enquiry@EuropeanResolverPolicy.Com for more details

The weakening of the resilience of the Internet, strengthening of the position of certain online platforms and online intermediary services into gatekeepers, and consolidation of user data to the detriment of privacy, all seem to be counter to the long-term interests of end-users. In addition, these steps appear to be counter to the recent Digital Services Act (and associated Digital Markets Act) initiative by the European Commission.

Greater transparency is desirable in the criteria used by Mozilla to approve resolver operators for inclusion within the TRR programme. Whilst the policy itself is published, it is unclear what process is applied to review and approve or reject participation. It is also unclear what commercial relationship, if any, is implemented between the parties.

Respecting privacy and security

We believe that privacy and security should never be optional on the Internet, and that as the developers of Firefox we have an important role to play in protecting our users from privacy and security risks. With that in mind, we have drafted our TRR policies with strict privacy requirements to minimize the potential that DNS data will be used for building user profiles.

> *At present the TRR policy does not refer to local legislation or regulations, implying that the TRR takes precedence over these. There has been some suggestion that this is to protect dissidents in undemocratic countries, however recent developments have shown that the regimes in these countries can take steps to block encrypted DNS. In such circumstances, there are tools better suited to the purpose than simply encrypting the DNS.*
>
> *In democratic countries, it is inappropriate for the policy of a tech company to subvert local legislation and regulation. Especially when doing so may weaken the privacy and security protections available to users, for example by substituting GDPR protections with the Cloud Act and FISA 702 f a resolver based in Europe is replaced with one operated by a US-company. Therefore the TRR policy needs to be amended reflect local requirements including those specified in legislation and regulations.*

We are interested in feedback on these privacy requirements, whether they can be tightened further, and what if any operational constraints they create.

1. Our current policy states that user data must not be retained for longer than 24 hours. A number of DNS providers, however, only keep data in ephemeral state and delete it almost immediately.

   Some jurisdictions may have legally binding data retention requirements which are likely to apply to resolver operators. In such jurisdictions, it should be possible for resolver operators to be able to participate in the TRR programme, therefore the terms ought to include an exemption to comply with legal or regulatory requirements. Allowance should also be made for data retention to support the functioning of any services that a user has opted to use, provided the data retention implications are made clear when the user activates those services.

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

   As noted above, the TRR should explicitly allow partners to retain DNS data in line with domestic privacy law and to optional features such as those related to customer service and cybersecurity.

2. What operational constraints, if any, are created by this maximum 24-hour retention time?

   Some optional services require data to be retained for longer periods to function. For example, some optional security and parental control services may need user data to be retained, both to operate the core service and also to provide certain functionality such as customised user reporting. Exceptions like these should not present a problem as the data retention requirements can be communicated to the user at the point that the service is activated, ensuring that their consent is obtained.

2. Are there exemptions that should be allowed by the policy for additional data collection in emergency circumstances? Please specify (e.g., the relevant circumstances as well as transparency and reporting requirements).

As noted in the opening comments, any interruption in a cloud-based DNS service, as has happened to, for example, Cloudflare's offering in the last few months, is likely to cause confusion, resulting in service calls to the user's ISP rather than the resolver operator. A curated offering like that provided through the Mozilla TRR ought to include an indication that the problem is caused by lack of availability of the DNS resolver together with provision for direct access to customer service in the event of any issues arising. The service provision should to accommodate both online and offline access, with support available in the local language(s).

3. Our existing agreements stipulate that providers in our TRR program shall undergo third-party audits to confirm compliance with our TRR policies; are there particular criteria (e.g., auditor qualifications) or considerations (e.g., cost) that we should take under advisement?

It is unclear what protections these audits will provide in regulated markets, assuming that the resolver operators comply with the local legislative and regulatory requirements. In so far as additional audits are deemed necessary, greater clarity should be given to the specification of the audit.

In regulated markets, consideration needs to be given to the audit cost to ensure that it does not serve as an unnecessary barrier to entry to smaller resolver operators, increasing the risk of centralisation whilst providing little or no user benefit.

4. Our current policy establishes that DoH resolvers in our program must maintain a transparency report providing public insight into the extent to which the resolver has been subject to government requests for data. How can this requirement be improved? What other mechanisms, processes, and governance tools may exist that could provide the public additional insight into such requests?

Any such obligations need to be consistent with local legislation and regulatory requirements.

Online safety

Numerous ISPs today provide opt-in filtering control services, and our deployment of DoH is designed to respect those controls where users have opted into them. We take very seriously the challenges presented by the breath of malicious, harmful, and illegal content present across the web today (indeed, Firefox uses Google's Safe Browsing service to protect Firefox users from malware and phishing websites). At the same time, we do not consider broad filtering and blocking through the DNS to be an appropriate means for ensuring online safety, since it entails significant risks to fundamental rights and is easily circumventable.

> *Whilst it is not complicated to bypass DNS filtering, most Internet users have no knowledge of the existence or function of the DNS. There is ample evidence in Europe that the vast majority of users do not bypass DNS filtering by opting to use a DNS provider other than their ISP (feedback from large European ISPs suggest that this is true for around 90% of consumers). In fact, the evidence suggests that those users rely on the network-based DNS filtering to keep them and their families safe, so anything that might weaken those protections needs to be approached with a great deal of caution.*
>
> *Security professionals have stated at various fora, including during discussions at the IETF and on IETF mailing lists, that it is unhelpful to remove a layer of protection. They have also asserted that multiple layers are beneficial, even if some are more effective at providing protection than others.*
>
> *There are applications-based security measures that can provide mirror the protections offered by network-based counterparts, although these do rely on users to keep them up to date, unlike the network-based services. Any protections built into a specific application such as a browser are much more limited in scope, only offering protection whilst browsing, unlike network-based services that can provide security and privacy protection for all Internet usage.*

With this in mind, we're interested in general feedback as to how **online safety goals can be met in ways that respect the technical architecture of the Internet and individuals' fundamental rights**.

More specifically, we welcome comments on the following technical questions related to online safety:

1. 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?

   The European Commission's Digital Services Act package, published on 15th December 2020, includes rules for online intermediary services such as ISPs and DNS service providers. Within the DSA package, online providers are required to do more to limit the spread of illegal content and goods, a requirement which does not appear to be reflected in the current TRR document.

A potential consequence of the current approach would be for some resolver operators to locate in jurisdictions with lax requirements. The resolver operator should meet the legal requirements that apply in jurisdictions where they are seeking to offer service (ie the primary country or countries of residence of their target users) and not just in the jurisdiction where the resolver operates.

Many markets have specific requirements concerning both content blocking and filtering that ought to be incorporated into any resolvers targeting users resident in those markets. Such requirements ought to be accommodated by resolvers serving users in those markets.

Local market requirements can include relatively straightforward requirements such as the implementation of content blocking to prevent access to copyright-infringing material and should be formal requirements backed, for example, by legislation, regulation or court orders. Some markets may have more specific requirements, such as the need for ISPs and mobile networks operated in the UK to block access to all adult content for all users unable to prove that they are aged 18 or over – this is usually enabled or disabled at the point of sale.

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

The argument that DNS-based filtering or blocking could affect fundamental rights is weakened considerably when the fundamental rights of those harmed by malicious, illegal or inappropriate content are taken into consideration. To consider three examples: (i) the distribution of child sexual abuse materials can be significantly reduced through the use of content blocking and filtering; and (ii) children can be shielded from adult content. In addition, malicious content can be either filtered before it reaches users or prevented from using DNS for its command and control functionality, reducing instances of fraudulent behaviour.

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

There are alternative methods to protect users from illegal and harmful content, with some services making use of SNI data to aid filtering. Deep Packet Inspection is also an option, although this is more intrusive as well as being more expensive to deploy and not scaling as well. Some tools use a combination of methods, for example, DNS filtering backed by selective use of DPI.

As noted elsewhere, solutions that are based on software installed on endpoints require action by users, both to install them and then to ensure that they are kept up to date. Network-based protections require no user action and may continue to provide protection to compromised endpoints.

4. 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.)

    1. 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?

    Visibility of the reason what content may have been blocked is helpful. The following is extracted from the European Resolver Policy[3] and should be added:

    > *"A description of the circumstances where an operator of a DNS resolver service MAY direct the user to alternative content and the nature of that content— for example to an explanatory web page whenever malicious content protection has been enabled and an attempt was made to look up a blocked domain name."*

    It would be beneficial if Firefox allowed the user to be redirected to a blocking page, for example due to a security policy or parental controls, providing it comes from the certified source of the trusted resolver. If example.com served in HTTPS is blocked for a security or parental control reason the DNS can respond to the resolution request with the IP of an explanatory blocking page. By doing this, the user benefits as the browser displays the reason for the block instead of a certificate error or blank screen.

    In addition, being clear how to challenge any incorrectly categorised and blocked or filtered content is important. The following is extracted from the European Resolver Policy and should be added:

    > *"Details of a complaints procedure should be provided to handle false positives and false negatives generated by any filtering or content blocking capabilities that are available."*

[3] See www.EuropeanResolverPolicy.Com or email Enquiry@EuropeanResolverPolicy.Com for more details

2. What challenges weigh against a requirement to publish block lists?

It is illegal in some countries to publish the location of illegal content. In addition, publishing a block list may make it simple to reverse engineer the blocking as well as (potentially) directing traffic to the blocked content, which may be harmful or malicious. It should also be borne in mind that publishing a block list provided by a third party may well infringe their copyright.

Noting these points, we recommend that the transparency requirements are more limited and instead require the resolve operators to publish their policy regarding blocking as well as providing details of any block lists and threat feeds that they may use for this purpose. This provides users with clarity about the scope of any blocking without falling foul of the problems outlined above.

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

Filtering provides clear benefits to users and care should be taken in the terminology used to explain it. Specifically, pejorative language such as "censorship" and "lying resolvers" is both unhelpful and inaccurate, betraying a lack of understanding of the needs and capabilities of the vast majority of Internet users.

The following is extracted from the European Resolver Policy[4] and should be added:

> *"An outline of any filtering options that are provided and details of how to opt-in/out of using these facilities. This information SHOULD NOT disclose information that would be helpful to those seeking to bypass or reverse engineer these filters."*

Care needs to be taken to ensure that, if a user opts to enable encrypted DNS within Firefox, they are aware of the potential impact in terms of greater exposure to malicious and/or illegal content, both of which could have direct, negative consequences for them.

[4] See www.EuropeanResolverPolicy.Com or email Enquiry@EuropeanResolverPolicy.Com for more details

Building a better ecosystem

Privacy and security issues differ across regions. As we seek to bring the protections of DoH to Firefox users in different regions, we're interested in general feedback as to **whether there are unique local considerations that we should be designing for in given jurisdictions.**

The current Mozilla TRR does not include any reference to the sharing of cyber intelligence. The following text is drawn from the European Resolver Policy[5] and should be added:

> *The resolver operator SHOULD share cyber intelligence information with appropriate stakeholders which may include national and regional Computer Security Incident Response Teams, cybersecurity agencies, law enforcement agencies, research institutions and other authenticated, benign third-party cybersecurity actors. Where cyber intelligence information is shared, it MUST first be anonymised*.*
>
> ** This has to be done using non-reversible anonymisation techniques that are consistent with the relevant rules and standards that protect users' personal data and privacy. See for example the Data Anonymisation Code of Practice from the UK Information Commissioner's Office.*

More specifically, we welcome comments on the following technical questions related to localisation:

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

   The current dialogue box that informs users that DoH is enabled (see below) is misleading, with no mention of the potential consequences of pressing the blue button. It is questionable whether selecting the "OK, Got It" button would constitute GDPR-level consent, especially as there is no attempt to explain what DNS is, even though very few users will have any knowledge of this.

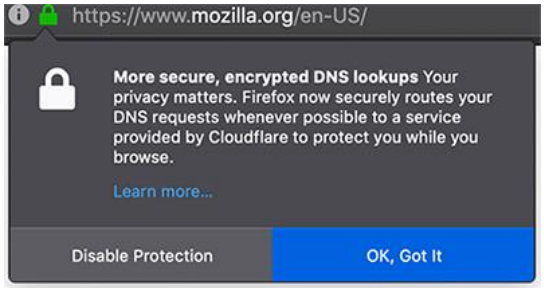

[5] See www.EuropeanResolverPolicy.Com or email Enquiry@EuropeanResolverPolicy.Com for more details

A better approach would be to provide a clear explanation of the possible consequences (for example regarding parental controls, malware protection, storage of personal data etc) so that the user is better equipped to make an informed choice.

A link to the resolver operator's transparency and privacy notice should be included, which in turn should include details about filtering and blocking policies together with clarify regarding the jurisdiction in which any data is stored and processed. In addition, a choice of resolvers from which the user can select their preferred option should be given.

One of the problems with the current TRR is that it does not refer to requirements such as GDPR or ePrivacy, nor does it take into account the recent Schrems II judgement from the European Court of Justice or the effect of US legislation such as the Cloud Act or FISA 702. A European version of the TRR would need to address these shortcomings if it is to be used within the EU or UK, noting that the current TRR model positions Mozilla as the Data Controller and the approved resolvers as Data Processors. One option to be GDPR compliant would be to adopt the European Resolver Policy referenced elsewhere in this document.

The current TRR text does not address the issue of data monetisation. It should specify that operators MUST NOT directly or indirectly monetise[1] any data arising from the use of these services[2] and SHOULD NOT enable other parties to do so either, without GDPR-level consent to do so.

[1] This is defined within the European Resolver Policy[6] as "Leverage for commercial or operational gain in any way. This includes but is not limited to: the sale of the data; machine learning based on it or associated anonymised data; leveraging the resolver operation in IPX peering deals; leveraging the resolver operation in the sale of CDN services to provide optimised performance to clients; other quid pro quo arrangements."

[2] This is defined within the European Resolver Policy as "This includes but is not limited to: Personal Data; IP addresses or other user or device identifiers; user query patterns consistently associated with a natural person or specific device from the DNS queries sent from the client; cache miss data."

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

Nil response.

[6] See www.EuropeanResolverPolicy.Com or email Enquiry@EuropeanResolverPolicy.Com for more details

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

Support for DoH amongst ISPs and DNS providers is likely to increase when a resolver discovery standard is agreed and implemented by client software from the major vendors. Such a discovery standard would need to support network implementations common outside of the USA, for example, the combination of DNS forwarders and private IP addresses (RFC 1918). More details can be found in https://datatracker.ietf.org/doc/draft-campling-operator-observations/.

At present Firefox is operating as a standalone application and is not taking into account any existing settings or user preferences on the endpoint. It is reliant on a curated list of approved resolvers, an approach likely to be overtaken when the work on resolver discovery that is underway within the IETF's ADD working group yields results.

It would be better if Firefox implemented the "same-provider auto-upgrade" (SPAU) approach that is used by other software including Chrome and Windows 10. Whilst neither has implemented SPAU in a manner that works with the network architecture commonplace in many markets outside of North America, it is nevertheless a step in the right direction. In the event that the user does not already have encrypted DNS in operation, it checks if the resolver operator offers an encrypted option and will automatically switch to this if it is available, maintaining user options and data privacy as it does not introduce a new data processor.

If support for SPAU is added, any user dialogue would need to be modified to acknowledge that encrypted DNS lookups are already in place before giving the user the option to decide whether they wish to change the existing, encrypted resolver for one that is part of the TRR programme (explaining the possible implications for doing so). If the SPAU resolver operator is part of the TRR programme then this step could be avoided completely. These points would also apply if an encrypted resolver has already been configured by the user or device owner.

More generally, being more transparent about the implications of using an encrypted resolver in terms of, in European markets at least, the GDPR impacts would be a positive step forward. As would support for DNS filtering given that it is widely used, making it simpler to enable encrypted DNS without placing additional requirements (and possibly costs) on the user such as the need to install, configure and maintain other software to replace these protections.

4. 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)?

Nil response.

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

Both Apple and Microsoft have added support for encrypted DNS at the operating system level since Firefox introduced its TRR programme. By deploying DoH in the browser, Firefox may bypass an existing encrypted resolver that has been set by the user or, in the case of an enterprise, by the device owner. Whilst the canary domain check is a good start, it would be better if Firefox was also able to determine if an encrypted resolver was already configured at the operating system level. In such circumstances, Firefox ought to use that resolver rather than connecting to a different one, or at least present this to an option to the user.

Additionally, as noted in the response to point 3 above, Firefox needs to support an SPAU option that functions on network implementations that common outside of North America, for example, the combination of DNS forwarders and private IP addresses (RFC 1918). More details can be found in https://datatracker.ietf.org/doc/draft-campling-operator-observations/.

Finally, as also noted in point 3 above, support for DNS filtering will avoid causing unnecessary operational challenges. Such filtering is widely used, making it simpler to enable encrypted DNS without placing additional requirements (and possibly costs) on the user such as the need to install, configure and maintain other software to replace these protections. An SPAU option would provide a straight forward way to achieve this.

How to respond

All responses should be submitted in the form of an accessible pdf or via email to the following address before 4 January 2021:

doh-comment-period-2020@mozilla.com

***NOTE: All genuine responses will be made available publicly on this Open Policy & Advocacy blog. If you wish for your submission to remain confidential, please explicitly indicate when submitting your comments by email.**

Submissions that violate our Community Participation Guidelines will not be published.