Vodafone's response to the Mozilla Corporation consultation "DNS over HTTPS (DoH) and Trusted Recursive Resolver (TRR)"[1]

20th January 2021

[1] Vodafone welcomes comments or questions on the views expressed in this document. They should be directed to Gianpaolo Scalone at [redacted]

Introduction

Vodafone is a leading telecommunications company in Europe and Africa. Our purpose is to "connect for a better future" and our expertise and scale gives us a unique opportunity to drive positive change for society. Our networks keep family, friends, businesses and governments connected and – as COVID-19 has clearly demonstrated – we play a vital role in keeping economies running and the functioning of critical sectors like education and healthcare.

Vodafone is the largest mobile and fixed network operator in Europe and a leading global IoT connectivity provider. Our M-Pesa technology platform in Africa enables over 45m people to benefit from access to mobile payments and financial services. We operate mobile and fixed networks in 21 countries and partner with mobile networks in 48 more. As of 30 September 2020, we had over 300m mobile customers, more than 27m fixed broadband customers, over 22m TV customers and we connected more than 112m IoT devices.

Given the scope and scale of our network operations, the security and privacy of the data of our customers is a major priority. We are therefore pleased to respond to this consultation being conducted by the Mozilla Corporation as we believe that the implementation of DNS over HTTPS (DoH) raises key policy issues that need to be addressed if the technology is to be deployed in a responsible way that does not weaken the security and privacy of our customers.

We believe that the encryption of DNS data 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. In particular, we are concerned that the technology may bypass network-based filtering that protect users from exposure to a malicious and illegal content, and that are used by parents to shield their families from age-inappropriate content. Further work is needed in this area by Mozilla.

We are concerned that any implementation does not restrict or reduce user choice, noting that the implementation of Mozilla's Trusted Recursive Resolver (TRR) policy in North America has resulted in a very limited set of three resolvers for users to select from, a significant reduction compared to those ordinarily available. This is an aspect of DoH deployment that Mozilla needs to give additional focus.

We also note that the implementation of DoH raises considerable communications challenges as it affects an Internet protocol, the Domain Name System or DNS, that is not known to, and certainly not understood by, the vast majority of users. Therefore any proposed deployments need to be carefully thought through so that they can be explained clearly to affected users. In our view, the current Mozilla approach has not addressed this issue.

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. This centralisation has negative consequences, in terms of competition and diversity at a key layer of the Internet stack, 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 [redacted] 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.

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, which is a flaw. For instance, Mozilla's specific set of privacy features and requirements seems to take precedent over the privacy laws and regulations that apply in the jurisdiction from which a user accesses Mozilla's services. The TRR does not explicitly commit to protecting user's domestically guaranteed privacy rights such as rights to redress or grounds for processing. We believe that Mozilla's TRR policy should explicitly require resolvers to respect the privacy law and regulations that apply to the jurisdiction from which the user accesses the services.

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 support 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?

For optional, over-the-top services, such as those linked to security and parental controls, it is not possible to delete user data within 24 hours as the data needs to be retained for a longer period both to operate the services and also to provide features such as user reports. Such exceptions should not be an issue however as the information retention policy for each service can be communicated to the user at the point of activation.

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 above, some optional OTT services require information retention beyond the 24 hour period. Given that the services are optional and that the information retention policy for each service can be communicated to the user at the point of activation, this should not be an issue but might be taken into TRR policy as an explicit exception to collect additional data if needed.

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?

We have no additional points to make.

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).*
>
> *Whilst application-based security measures have their place, they are reliant on users keeping the software up to date, unlike network-based services. Browser-based protections are of course limited in scope, only protecting the user whilst browsing, whereas network-based services can provide security and privacy protection for all Internet usage and are updated by the ISP without needing to involve the end-users.*
>
> *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. It should be noted that DNS filtering provided protection against the recent SolarWinds attack.*

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. As required within the DSA, online intermediaries will have to do more to limit the spread of illegal content and goods.

A potential consequence of the current approach within the Mozilla TRR 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.

Looking at an example in a specific market, it is a requirement in the UK for all ISPs and mobile network operators to block adult content to all users unable to prove that they are aged 18 or over. This is typically enabled or disabled at the point of sale, although the settings can subsequently be changed by the user providing that have appropriate proof of their age. A cloud-based resolver operator targeting UK users should implement adult content filtering for those users. The same applies to other EU markets, for example Italy.

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 or illegal content are taken into consideration. For example, the distribution of child sexual abuse materials can be significantly reduced through the use of content blocking and filtering, and children can be shielded from age-inappropriate content. In addition, malicious content can be filtered before it reaches users, reducing instances of fraudulent behaviour and other activities that may be harmful or weaken user privacy.

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 it doesn't scale as well. There are some tools that use a combination of methods, for example, DNS filtering backed by selective use of DPI.

   The problem with solutions that require installation of software on endpoints is that they require user action. Network-based protections can block access to illegal and harmful content by compromised end-points and by users able to circumvent, disable or not even install local controls.

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.)

   A simple list of blocked domains by country is problematic as it effectively provides a directory of malicious and illegal content. In any event, the publication of such a list is likely to be illegal in some jurisdictions. More details are included in response to the further questions below.

   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, if 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.

      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 [redacted] 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 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, whether that's in terms of protecting them from malicious content or, in the case of children, protecting them from age-inappropriate content. Care needs to be taken to ensure that the implementation of encrypted DNS by Firefox doesn't bypass such filtering to the detriment of the security and privacy of the user.

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."*

See also the comment in the response to 4.1 above regarding the benefit of redirecting to a blocking page, for example due to a security policy or parental controls, if it comes from the certified source of the trusted resolver.

[4] See www.EuropeanResolverPolicy.Com or email [redacted] 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 in Firefox that informs users that DoH is enabled (see below) is misleading, with no mention of the potential consequences of pressing the blue button. For example, there is no mention that opting-in may bypass network-based protections, nor is there any indication that it may involve a change of data processor and/or jurisdiction in which personal data is stored.

   By presenting a single resolver option, Firefox seems to be operating as a data controller. The dialogue box does not appear to meet EU or UK GDPR-level consent requirements.

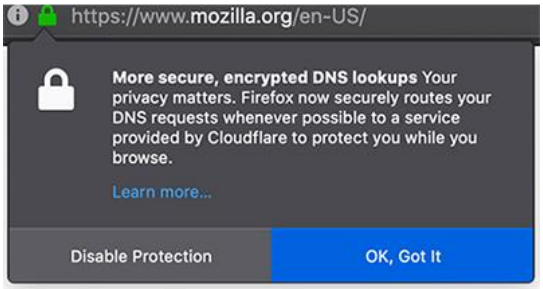

https://www.mozilla.org/en-US/

More secure, encrypted DNS lookups Your privacy matters. Firefox now securely routes your DNS requests whenever possible to a service provided by Cloudflare to protect you while you browse.

[5] See www.EuropeanResolverPolicy.Com or email [redacted] 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. It should provide a link to the resolver operator's transparency and privacy notice, which should include details about filtering and blocking policies and, ideally, more than one resolver from which the user can select their preferred option.

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, with one option being to comply with the European Resolver Policy.

In addition, 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 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?

   We have no additional points to make.

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

Support for DoH (and other encrypted DNS standards including DoT and DoQ) 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/.

Currently, Firefox is operating as a standalone application and is not taking into account the settings on the host device. As noted in response to the next question, it would be better if the browser was 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.

What is also lacking in the current Firefox software is the so-called "same-provider auto-upgrade" (SPAU) option that is already implemented in the Chrome browser and Windows 10 operating system, albeit with some limitations. By implementing SPAU, Firefox would detect that the current unencrypted resolver operator also has an equivalent encrypted offering and automatically migrate to that without the need for any user intervention as there are no changes to any services or options.

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)?

We have no additional points to make.

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)?

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, device owner. It would be better if, in addition to the existing checks, Firefox was 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.

In addition, as noted in the response to point 3 above, Firefox needs to support an SPAU option that functions on 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/.

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:

[redacted]

***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.