Stay Open. OX®

Cologne, 20 January 2021

Open-Xchange's contribution to Mozilla's online consultation on the DNS-over-HTTPS and Trusted Recursive Resolver policies

Open-Xchange, as the leading European open source software company in the email and DNS space, and as the service provider of choice for some of the biggest Internet access providers in Europe, would like to thank Mozilla for the opportunity to provide comments on policies for the rollout of DNS-over-HTTPS and for the selection of *"Trusted Recursive Resolvers"*.

As a free software company, **we appreciate Mozilla's openness and commitment to a free and privacy-friendly Internet**, as we share similar values; in fact, in many other discussions and campaigns we find ourselves on the same side of Mozilla. We also **share the objective of a prompt and widespread deployment of encrypted DNS protocols** to increase the user's security and privacy.

At the same time, **we often expressed concerns on Mozilla's unique approach to encrypted DNS deployment** and we will take this chance to explain them in full.

General considerations

Before getting into each of the sections of the questionnaire, we would like to start by addressing two general issues that underpin Mozilla's strategy as a whole.

#1 – Per-application DNS resolution

Immediately since we learned of it, **we took issue with Mozilla's baseline assumption that it is the proper role of a browser to determine, restrict, or even suggest the use of a specific DNS resolver**, different than the one configured in the operating system, to its users.

Moving the choice of DNS resolver from a device-wide setting to a per-application setting reduces the user's ability to control where their personal information is going. Up to now, conscious users that want to keep control over their DNS resolution just had to set their DNS resolver once for all in their system configuration; in Mozilla's new paradigm, they now need to remember to set their DNS resolver in each and every application. **Per-application DNS resolution could perhaps be an option that advanced users actively turn on to deal with specific situations**, but not the default behaviour recommended to average users.

Also, should this paradigm become widely accepted, we would possibly see some less honest and user-focused applications start to direct the user's DNS and other data to services of their liking, for monetization and tracking purposes. Conversely, users that already employ DNS-based systems (such as PiHole) to protect their privacy and filter out tracking advertising would find themselves exposed to privacy leaks if the browser moved to a different resolver, sometimes without even realizing that this is happening. **Keeping the users in direct and easy control of their network settings is a key requirement to protect their privacy**.

Additionally, **if different applications started to use different DNS resolvers, it is very likely that they would sometimes get different answers to the same queries**, for many possible reasons (not just security and filtering policies, but also different caching regimes and timings, or even just differences in implementation and functioning). This would lead to utterly confusing user experiences, in which for example two browsers pointed at the same website at the same time show different things.

Mozilla does not seem to have considered in full how to deal with these situations, and especially how to provide support to users in case of problems. Until now, DNS failure is considered as a form of network failure, and the Internet access provider will take care of supporting the user. **Browsers that adopt a different resolver than the system or network one should make sure that they also provide adequate error detection and support options for average users** in case their own resolver malfunctions.

#2 – Appropriateness of DNS filtering for online safety

We acknowledge Mozilla's opinion, as stated on page 3 of the questionnaire, that *"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"*.

However, since Mozilla asked for comments, **we would like to challenge that opinion on multiple levels**.

First of all, taken literally, Mozilla seems to think that that the problem does not lie in blocking in itself, but in the fact of using DNS to that purpose. We know well all the limitations of DNS filtering, but we note that **it emerged during the last 20+ years as the filtering and redirecting mechanism of choice in multiple uncorrelated use cases**, from law-sanctioned blocking of illegal content to real-time disruption of botnet command and control mechanisms, from corporate *"split horizon"* intranets to removal of tracking advertising, from opt-in parental controls for children to network health monitoring and advance threat detection by ISPs and corporate network administrators. **If DNS filtering for online safety is so common, it is indeed a sign that it is not such a bad mechanism after all**.

Specifically, **the two reasons provided by Mozilla to justify its assessment seem to contradict themselves.** First of all, when discussing DNS filtering, Mozilla seems to consider just the use case of law-sanctioned content blocking, which, we stress, is just one of many, and is far from forming the majority of the uses. Moreover, for that specific case, **DNS blocking has emerged exactly as the mechanism that reduces the risks to fundamental rights by being easily circumventable**; and this has been the conscious reason why, in democratic countries where law-mandated content blocking is regularly in use, local Internet communities have been pushing for DNS as the implementation vehicle.

Other alternative mechanisms, such as IP-address-based blocking or deep packet inspection, are much more prone to over-blocking, cannot be so easily bypassed, and prompt much deeper intrusions into the user's privacy. Indeed, **by making DNS blocking unavailable, the result in many countries and in many situations could be the move to broader and more intrusive blocks**, and not the opposite.

On the other hand, we wonder whether what Mozilla actually means is that blocking access to content at the network level is always inappropriate, independently from whether it is effected via DNS or via other mechanisms. If so, perhaps Mozilla could state this more clearly, as it would avoid discussing the wrong set of questions.

We understand and share Mozilla's concerns about free expression online. However, **we note that free expression is not the only fundamental right that needs to be protected**; in current times, we are in the middle of broad waves of hate speech and racism, organized attacks against minorities, foreign political propaganda, terrorist recruitment, child sexual abuse material and other types of harmful content that endanger the wellbeing, the dignity and the life of actual people and democracy in general. We also see counterfeit drug shops, Covid-19 related scams and anti-vax propaganda, unlicensed gambling, all types of phishing and malware distribution websites, and other types of content that can seriously damage users that unknowingly click and end up on them. **Blocking access to harmful content is not always the right answer to the problem, but in many situations can be an effective mechanism – often, the only practically and promptly available** – to protect the average user from the damage that such content could create.

Importantly, **the decision on whether the free speech rights of users posting such content should prevail over the rights of the users that such content would damage is not an easy and clear-cut one**; it needs to be subject to careful evaluation that depends on the specific situation and on the legal framework under which it falls. Assuming that content filtering is never appropriate trumps any chance of such careful evaluation, forcing a one-size-fits-all approach that will do more harm than good.

We also note that almost all online safety use cases – from parental controls to botnet stopping, way beyond free expression issues – need to deal with an adversarial endpoint: a child that may try to bypass the family blocks, or a bot application trying to reach its control center by any means. **Endpoint-based content control mechanisms are only useful in a limited subset of the use cases where in-network blocking is adopted today**.

In conclusion, while we appreciate free speech concerns, **we think that Mozilla has missed the need for a proper, complex, nationally variable balance between free speech and other rights, and has way underestimated the relevance of DNS blocking as an instrument for other purposes**, first of all network security, and the damage to the reliable and safe functioning of the Internet that would be determined by the disruption of all DNS-based blocks.

Our suggestion, in the end, is that **Mozilla should revise its position on this matter and come to a more nuanced policy stance that acknowledges the usefulness of DNS filtering** in a broad number of use cases.

In any case, if Mozilla believes that other technical mechanisms for blocking access to unsafe online destinations are more appropriate, we stand ready to hear Mozilla's practicable suggestions on how to deal with the many use cases currently addressed via DNS-based blocks.

Finally, even if Mozilla decided to stay with its current position, **we challenge the assumption that it is appropriate for Mozilla to make a judgement on the matter on its own for the entire Internet community** and push it onto everyone else. We note that Firefox adopters do not use the Internet in isolation, so the stakeholders in this matter do not just include those who develop and use the browser, but also those that provide and manage the networks, servers and services used through the browser.

As discussed above, the existence and extent of DNS blocks affects multiple stakeholders – users, ISPs, governments, and groups that may be harmed by the easy accessibility of hateful and threatening online content, ranging from children to endangered minorities. **Any conclusion on the appropriate balance between blocks and access rights should be reached only through a multi-stakeholder discussion in the appropriate democratic venues**; such conclusion is likely to vary across different countries, as situations, histories, values and legal customs vary.

The evaluation of whether any content control mechanism *"entails significant risks to fundamental rights"* lies within the purview of each country's Parliament, Constitution and related courts, which indeed, in a number of European countries, expressed opinions on that very issue and have worked or are working on specific regulations. The proper balancing of rights cannot be unilaterally decided by a browser maker for all countries and cases.

We thus suggest that Firefox, and all browsers in general, should consider themselves neutral platforms that connect users with websites and online content, and should **focus on being flexible enough to support any national multi-stakeholder agreement on content control achieved through the established democratic processes**. Mozilla is welcome, if so desires, to join local Internet communities in these discussions and advocate its own position, but it should not use its technical role to push its own views over the local democratic consensus.

General comments regarding Mozilla's TRR policies

As explained above, **we think that the very idea of a per-browser *"Trusted Recursive Resolver"* is not a good one**. However, should Mozilla choose to continue with the concept, we provide some specific comments on the TRR selection policy.

We note that, in Europe and other parts of the world, strong privacy laws already provide significant guarantees for the user's privacy – the need is more for adequate law enforcement than for additional rules by browser operators. However, specifying the best agreed technical implementation of privacy law principles in DNS operations can be useful.

We are very concerned by the idea that each browser, operating system or device maker could come up with their own resolver accreditation policies that DNS resolver operators needed to meet before being allowed as a recommended choice, or even as a choice at all, within the maker's applications and devices.

Multiple sets of possibly incompatible requirements would make it hard and costly, or even plainly impossible, for DNS operators to be accredited everywhere, especially for smaller operators and non-profits. Even smarter users running their own private resolvers could be affected by this, if accreditation became a requirement for use.

Also, **unhealthy market practices could emerge**, with DNS operators compensating browsers for accreditation or for becoming their default or exclusive DNS resolver, something that would alter competition in the DNS market and could even be illegal under current and upcoming Internet anti-trust regulations. This could also create an inherent push for DNS operators to recover their costs by finding ways of monetizing the data they get from the DNS operation, circumventing the privacy conditions of their accreditation as far as possible.

The concept of per-application DNS resolver accreditation mechanisms is thus dangerous, as it would lead to further centralization; **only a relatively small number of bigger operators could be accredited and thus recommended, or even available, to users**. The higher the barriers for accreditation, the stronger the push to centralization. We note that the current list of accredited resolvers in Firefox is extremely short; moreover, it only includes commercial DNS operators, but not non-profit, privacy-oriented public resolvers (e.g. Quad9), something that seems to clash with the very objective of the initiative.

As a minimum, **efforts should be made to establish a single industry-wide resolver accreditation policy**, so that operators need to meet just one set of requirements and perform just one accreditation process. These requirements and processes should be entirely or almost free of charge and should be designed not to be exclusionary against smaller independent operators.

Regarding the content of Mozilla's current TRR accreditation policy, we thank Mozilla for clarifying that abiding by applicable national laws in terms of content blocking and law enforcement transparency should not be a cause of exclusion. We note however that **clauses #2 and #3 in the *"Blocking & Modification Prohibitions"* section may actually conflict with national laws in a number of countries**, as often there are requirements to keep lists of blocked domains confidential (for example for domains related to child sexual abuse, for obvious reasons), or to redirect users trying to access a blocked domain to a specific legal message. Exceptions for national rules should be added to these points as well; given the current regulatory trend in Europe, which requires non-European Internet services to apply European laws when operating with European users, we note that these may not just be the national rules of the operator's country, but also the national rules of each user's country.

On clause #3 specifically, we share the objective to prevent responses for blocked or non-existent domains to include advertising or tracking – something that however, at this point in time, is quite uncommon in Europe. However, in the case of blocked domains, it would be a much better user experience – rather than just returning a NXDOMAIN response – if the user could be safely redirected to an error page explaining that the domain is blocked and why, also providing options for configuration changes or recourse where appropriate. Standardization work on this is ongoing at the IETF (e.g. see the specific extended error codes included in RFC 8914). We suggest that **Mozilla should participate in this work and make such a practice compatible with its TRR policies**.

Finally, we are concerned at the principle level by Mozilla's stance that resolvers may still be removed from accreditation for complying with national content blocking laws, *"depending on the scope and nature of those obligations"*. This equates to Mozilla claiming the role of judge of adequacy and democracy for the national laws of each country, which is highly inappropriate and, for democratic countries like the European Union's, almost offensive. Also given recent events in the United States, we note an increased concern among European political leaders about global Internet companies unilaterally assuming the role of judges in free speech, and we dare to suggest that Mozilla should refrain from positioning itself in the middle of that controversy.

Respecting privacy and security

We note that **Mozilla seems to have a rather different understanding of privacy than what is common in Europe**. Mozilla's claim is that *"we have an important role to play in protecting our users from privacy and security risks"*, implying that it is the proper role of browser developers to determine which parties are to receive the user's data and at which conditions, including adding new parties (such as Cloudflare or other TRRs) to the data loop, independently from whether the users are actually and fully aware of what is happening.

On the other hand, **the basic concept of privacy in the GDPR and in other European regulation is that the user is the only party that should control who gets to receive their data**, and that **one less party in the data loop is always better than one more**. Any party wishing to help users protect their privacy should rather inform them in detail, provide choice and acknowledge their final decisions.

In terms of data retention, we note that the feeding of DNS query details into statistical engines and machine learning algorithms is already of great economic value even if those query details were then not logged at all. Potentially, even politically sensitive information, such as the readership profile and trends of specific online websites, could be acquired in this way. While it is important to reduce the duration of any data retention policy, **it is more important to control what the resolver operator does with the data it receives**, independently from whether and for how long these data are stored in logs – while, on the other hand, some extent of logging is necessary for technical purposes (e.g. debugging, performance optimization).

In this regard, while we find the list of prohibitions in the TRR policy quite exhaustive, we note that the prohibition to store *"non-aggregate anonymized data"* and the prohibition of combining data *"to identify individual end-users"* seem to leave open the possibility for the operator to build and maintain permanent profiles of DNS users, even combining DNS data with other data sources, as long as these profiles are not associated with an actual identity, and as long as the data that feed the profiling algorithm are discarded after use. If this is not Mozilla's intention, perhaps **there could be a clearer prohibition of any kind of user profiling, even if anonymized**.

Regarding third party audits, we note that **unless these audits are somehow permanent, it will always be possible for the TRR operator to change their data retention policies silently just after the audit ends**. The attempt to audit TRR operators is commendable, but in the end, there is a fundamental and unavoidable need for trust, which is way more important to determine whether TRR commitments can be realistically maintained.

In this regard, we suggest that **Mozilla should also acquire and consider the motivations that lead an entity to deploy and maintain a TRR-compliant public resolver on a consumer scale**, something which can entail very significant costs. Those motivations are ultimately going to determine whether the service will survive and function reliably in the long term, but also whether internal revenue pressure may determine open or stealth attempts to exploit and monetize the data beyond what is publicly stated.

As stated in the previous section, we note that **any transparency requirements on law enforcement access** – a concept which seems to be particularly relevant only in the United States – **are likely to conflict with existing national laws** that mandate the confidentiality of police investigations, which are in force in most of Europe (possibly everywhere). Besides this, we think that it is the duty of every good citizen to cooperate with police investigations, and while we can imagine abusive scenarios of politically motivated investigations on specific people, none of them seems to be realistically possible in the European Union.

Online safety

As explained in the general considerations, we challenge Mozilla's assumption that it is their role to judge whether DNS blocking and filtering is *"an appropriate means for ensuring online safety"*. This judgement is to be made by governments and Parliaments, for what regards national content blocking laws; by ISPs, for what regards the security of their network; and by users, for what regards the online safety of themselves and of their family members. We think that **the browser should respect the decisions made by these stakeholders and avoid disrupting any DNS-based safety mechanism that other stakeholders want to adopt**.

As developers of open DNS resolver software, we are asked by several of our users and customers to add filtering capabilities, because of the legal requirements that impose the filtering of specific destinations, or because the users of our software want to turn on filters to increase safety for their own end-users. It is up to our users to determine if, when and how to use the DNS filtering capabilities; even if we wanted, as PowerDNS and dnsdist are pieces of free software, we would not even be able to prevent anyone from using and modifying them for DNS filtering purposes – after all, this is part of the very idea of free software.

However, we are particularly interested in the outcomes of questions #2, #3 and #5 in this section, as **we would be happy to work with Mozilla and with any other interested party to address any issues that derive from the use of DNS filtering**, both at the technical and policy level. In the past, the refusal – even in some technical standardization venues – to have a practical discussion on these matters has led to a lack of standardization and thus of transparency and inclusivity, with each ISP implementing their own software, standards and procedures. This disadvantages small operators that often are subject to the same filtering requirements as any other (bigger) ISP, but without the same capabilities to implement these filters well, when free software solutions, standard procedures and open data flows are not readily available.

For example, **we share the concerns about blocking practices that do not include an easily available mechanism for recourse against any possible mistake or abuse**. In our experience with telco customers in countries (like the United Kingdom) where DNS filters are common, ISPs are more than willing to act promptly whenever they are noticed that their filters are not working properly or are affecting the wrong domains; to the best of our knowledge, this issue is not as broad as it is sometimes depicted. Again, we submit that an implementation of advanced error pages, building on RFC 8914, would do a lot to address this problem in practice, as it could include an interactive way to submit a warning about a mistake.

Regarding question #4, we again stress that **full transparency may be impossible for specific types of filtering**, where publishing lists of blocked domains may be prohibited by national laws or, more importantly, could harm the people that filtering aims to protect; for example, the release of a list of domains blocked for distributing child sexual abuse material would practically create a ready-to-use directory for child abusers. In these cases, transparency could be achieved by adding trusted third parties (e.g. digital rights NGOs) into the data loop, so that they could verify how the list is being used while keeping it confidential.

Moreover, when blocking access to domain names, time is often of the essence. Real harm may be done by deferring the block even by a few hours; for example, phishing campaigns often go out all of a sudden and aim to catch all their victims in the first few hours or so; botnets use domain generation algorithms that change the controlling domain name at regular intervals that can be very small; and so on. **Any transparency and validation procedure should be designed not to introduce delays in this kind of blocking**.

At the same time, **the standardization of formats and mechanisms for compiling and transmitting domain blocklists would also naturally increase transparency and accountability and reduce opportunities for abuse**. For example, lists could be cryptographically signed by the appropriate public authority or abuse list provider and distributed securely to any party that needs to implement them. Again, we would be happy to contribute to community work on this topic.

Building a better ecosystem

We thank Mozilla for the intention of *"building a better ecosystem"* in the DNS and ISP industry. However, we suggest that a company like Mozilla, never having been part of such ecosystem in the first place, should learn in detail how the current ecosystem actually works before trying to change it for the better. Indeed, we trust that the discussions of the last three years, no matter how controversial, helped Mozilla in understanding that things in the DNS world are way more complex than they look at first sight, technically, legally and policy-wise.

So, to answer question #3, we think that **the best way to gain support of the DoH standard among ISPs and DNS providers is to engage with them and accommodate their views and needs**.

More specifically, **we support the so-called "same provider auto-upgrade" (SPAU) deployment model**, in which the client application tries to determine whether an equivalent encrypted resolver exists for the unencrypted resolver currently configured in the system and, only in that case, upgrades the DNS connection automatically to DoH. We think that it would be preferrable if browsers would just ask the operating system for a system-wide DoH service if available, but in case clients still want to adopt the per-application resolver selection model, the SPAU principle avoids most of the issues that derive from a unilateral change of resolver operator made by the browser, as per Mozilla's current deployment model.

Other major players (Google, Microsoft) have adopted the SPAU model, and it would be highly beneficial to everyone, and to DoH deployment, if it became the industry standard. Moreover, **the SPAU model encourages the parties that are currently providing consumer resolution services to provide a DoH service as well**, so that the user does not switch to another provider's encrypted resolver instead. In other words, the SPAU model is the natural answer to Mozilla's question #3.

On question #4, we agree that **when roaming onto unknown public networks, such as those of hotels and public spaces, DoH could provide additional privacy and security**; its deployment in those situations is particularly welcome.

At the same time, we note that even in that case DoH is largely insufficient to grant full privacy. When real risks are at stake, such as for the oft-mentioned case (but not applicable in Europe) of a political dissident in an authoritarian country, we think that people should be led to use trusted VPNs, obfuscated routing via Tor, and other mechanisms; they should not think that encrypting their DNS connection is enough. Implementations that detect DoH traffic without decrypting it already exist.

Additionally, we note that **there are also many cases in which computers rarely or never roam**. Fixed computers of home and corporate users never leave their local network; corporate laptops are often meant to stay on their own corporate VPN, including for DNS resolution. **It cannot be assumed that each copy of Firefox will be running on a computer that roams** or that will ever be other than on a local network with a trusted local resolver provided as part of its Internet access arrangements.

Finally, we remind Mozilla that **there is a small but strong minority of technical users that takes the resolution of their DNS queries directly into their hands**, by carefully choosing their DNS provider or even by deploying their own DNS resolver on their own private server, often with custom DNS filtering capabilities (see for example the aforementioned PiHole project). DoH can help these users by making sure that their queries cannot be intercepted or mangled in-transit.

While running one's own resolver may have negative performance implications and may give up the additional anonymization deriving from merging the queries with those of other users, if done properly it can be a way to meet the privacy objectives that underlie DoH. Such a DNS resolution model is not adequate for most users, but for those who can, it would be perfectly in line with the original Internet principle of decentralization, much more than the idea of *"Trusted Recursive Resolvers"*. **We thus suggest that Mozilla keeps the *"personal DNS resolver"* use case in mind when designing its DoH approach**, not bypassing such a server when it has been configured at the system-wide level.

In the end, **any DoH deployment model chosen by Mozilla should take into account all common use cases and situations**, rather than only some of them, and make sure it can work in each of them. Failure to do so creates pushback against the adoption of DoH as a whole, as we could clearly see in the last few years. Dismissing many use cases as inappropriate or wrong, only because they are other people's use cases and they do not fit within one's own initial assessment of the landscape, is not going to *"be disruptive"* and *"create a better ecosystem"* – it is just going to create problems for everyone.

We thank Mozilla for the opportunity to provide our exhaustive, sincere and direct comments on this matter. Again, no offense is meant by them and we reiterate our general appreciation for one of the most successful free software projects ever. We hope to continue this discussion and to work together on DoH implementation even more than we did in the past.

About Open-Xchange

Open-Xchange (OX) is a developer of secure and open communication and office productivity software, IMAP server software and DNS solutions. Since 2005, it has partnered with many of the world's largest Internet Service Providers (ISPs), telcos and carriers to deliver open source email and productivity solutions that include secure storage, file and document management. OX Dovecot is the world's most popular IMAP server software and OX PowerDNS provides secure DNS services to telco customers and their users worldwide. Software developed by Open-Xchange is used by 200 million people globally. It is headquartered in Germany, with offices in Bremen, Cologne, Dortmund, Hamburg, Nuremberg and Olpe, and international offices in Finland, Italy, Japan, the Netherlands and the USA.

Contact

For any enquiry about this document please refer to Vittorio Bertola, Head of Policy & Innovation at Open-Xchange, [redacted] *.*