January 20, 2021

VIA EMAIL: [redacted]

Mozilla Corporation
331 East Evelyn Avenue
Mountain View, CA 94041

Re: Mozilla's Questions for Comment Regarding DNS-Over-HTTPS Implementation

Dear Mozilla:

Access Now appreciates the opportunity to submit comments in response to Mozilla's *Comment Period on DNS-over-HTTPS (DoH) Implementation*.[1] We support Mozilla's efforts to improve privacy, security, and online safety through browser-based implementation of DoH. We provide the following responses to some questions and statements made in the questions for comment.

I. Respecting privacy and security

The right to privacy is a fundamental human right, and should be respected globally. It is important that privacy protections take into account context. We provide the following comments to ensure context is respected.

First, DoH should be on by default, but should not be the only option for people who use Firefox. Second, the 24-hour window should not be lengthened, and could even be shortened. Third, third-party audits should be mandatory. Fourth, Trusted Recursive Resolver (TRR) transparency requirements should be improved.

A. DoH on by default

In Mozilla's questions for comment, Mozilla states "[w]e believe that privacy and security should never be optional on the Internet." If by this statement Mozilla means to force all Firefox users to encrypt their DNS traffic over DoH, we would oppose this proposal. There may be legitimate reasons why a person may not want to do so. For example, users seeking remote malware analysis may face difficulties with that analysis being restricted if DNS traffic is forced to DoH in all circumstances. Other examples include the challenges that enterprise system administrators may face when it comes to custom landing pages and other organisation specific policies.

[1] Questions for Comment, Mozilla, https://blog.mozilla.org/netpolicy/files/2020/11/DoH-Public-Comment-Period-Question-for-Comment.pdf.pdf ("QFC").

DoH should be the default option, as it currently appears to be in Firefox browsers, but not the only option. Mozilla should also allow people to choose which TRR they prefer, as the browser currently allows. Such practices benefit people because, while most prefer not to change their default options, those who desire to make the change would be at liberty to do so.

B. The 24-hour data retention policy

Mozilla asks "[t]o what extent can our [24-hour data retention policy] be shortened further while allowing providers sufficient data to operate the service?" Data retention requirements should always be set at "only as long as is necessary to provide the service" to minimize potential privacy intrusions caused by, for instance, companies engaging in detailed, creepy surveillance of people that use their service or unauthorized access to and disclosure of data. In this instance, the limit should be as short as is necessary for the TRR provider to provide the service.

From our perspective, under no circumstances should the 24 hour limit be lengthened, and there is a strong argument for it to be shortened. As Mozilla states, many DNS providers delete data almost immediately, which undercuts the need for TRR providers to retain it for longer. Further, most DNS issues become quickly identified and should generally be obvious to service subscribers and monitoring systems.

Mozilla could set a rule that TRR providers will delete data as soon as it is no longer necessary to provide the TRR service, and should be deleted no later than 24 hours after it was collected. It could further state that in a certain amount of time, perhaps six months, after implementing the policy, it would check with TRR providers to review implementation and to see if the timeframe should be adjusted downward.

The current policy states that "[o]nly aggregate data that does not identify individual users or requests may be retained beyond 24 hours." Mozilla should define "aggregate data" narrowly. Aggregate data should be defined similar to telecommunications privacy law, which is straightforward: "collective data that relates to a group or category of services or customers, from which individual customer identities and characteristics have been removed."[2] Thus, aggregate data is *both* (1) data relating to groups or categories of services or customers, and (2) de-identified. As you likely know, merely de-identified data can be re-identified in many situations, and browser history is one of the more easily re-identifiable datasets.[3] In addition to the more common examples of the NYC Taxi Commission and Netflix movie debacles,[4] a recent paper showed that machine learning can re-identify data that has been de-identified, and that

[2] 47 U.S.C. §222(h)(2).
[3] Luc Rocher et al., *Estimating the success of re-identifications in incomplete datasets using generative models*, Nature Communications (2019), https://www.nature.com/articles/s41467-019-10933-3.pdf ("In 2016, journalists re-identified politicians in an anonymized browsing history dataset of 3 million German citizens, uncovering their medical information and their sexual preferences.").
[4] Rajesh Parthasarathy, *Reidentification Risk of Masked Datasets: Part I*, Forbes (June 9, 2020), https://www.forbes.com/sites/forbestechcouncil/2020/06/09/reidentification-risk-of-masked-datasets-part-i.

99.98% of Americans can be reidentified with 15 data points.[5] Thus, Mozilla should clearly state that TRR providers may not retain merely de-identified datasets; the data must meet the definition of aggregate stated above.

C. The third-party audit requirement

Mozilla next asks questions regarding their third-party audit requirements. The third-party audit of TRRs is vital to the success of the DoH system and should be one of Mozilla's primary focuses. Requirements of third party audit should not be altered or waived because it satisfies a key need for 'proof of work' for TRRs to demonstrate trust in network activity and combat malicious cyber activity.

Mozilla should seek to further develop avenues that would ease the burden on civil society organizations that wish to participate as TRRs. As a civil society organisation working on digital security and embedded within a larger community of civil society and public interest technology actors, the cost incurred by an audit requirement may deter civil society involvement in the TRR space to the detriment of the program. Civil society organizations working on digital security and public interest technology may wish to assist or participate as a TRR, building on past successful efforts in the SSL/TLS and open certificate authorities areas. Mozilla could, for instance, seek trusted third-party auditors in a position to provide pro bono audits to civil society digital security organizations, or attempt to find funding resources for potential civil society TRR initiatives, or at least seek to allow for mechanisms by which public interest funders and digital security assistance groups could help cover the cost of such audit requirements for civil society.

Enforcement against malfeasors may be necessary. Mozilla should have in its policy an enforcement section stating Mozilla reserves the right to disqualify any TRR, TRR provider, or auditor not following the guidelines or enforcing the policy in a good faith manner.

D. Transparency requirements for TRRs

The request for comments state that "[o]ur current policy establishes that DoH resolvers in our program must maintain a transparency report providing [information about] government requests for data." It then asks "[h]ow can this requirement be improved?" Given our extensive experience with transparency reporting,[6] we have several ideas for how the policy can be improved.

The current policy is straightforward but incomplete: "[t]here must be a transparency report published at least yearly that documents the policy for how the party operating the resolver will handle law enforcement requests for user data and that documents the types and number of requests received and answered, except to the extent such disclosure is prohibited by law." This

[5] Rocher, *supra*, note 3.
[6] Isedua Oribhabor & Peter Micek, *The What, Why, and Who of Transparency Reporting*, Access Now (Apr. 2, 2020), https://www.accessnow.org/the-what-why-and-who-of-transparency-reporting.

policy keeps a lot of publicly beneficial data private when there is no reason to do so. The policy can be improved by

- requiring reporting on all types of government requests, not only law enforcement;
- requiring reporting on actions the TRR provider has taken on user accounts in enforcement of its own terms of service;
- breaking down stats on requests by country; providing stats on requests rejected or challenged and why;
- stating policies on notification to users when their data or accounts have been requested by government authorities;
- stating policies on any form of remedy or appeals process available to users;
- ensuring accessibility of transparency reports that are both human readable (available in local languages where applicable), easy to find on the resolver's website (a dedicated webpage or address), and available in a machine readable format so that the reports can be collated automatically and compared across the world; and
- indicating that transparency reports should ideally be published twice a year, with a minimum requirement of at least once a year.

In addition to a transparency report, the TRR provider should be required to adopt a human rights policy stating its commitment to respecting its users' privacy and other human rights, and outlining the steps the company will take at the senior management level to ensure that the policy is enforced throughout the TRR's business operations.

II. Online safety

Mozilla asks a series of questions about online safety in its request for comments. Below, we argue first that blocking and filtering at the DNS level presents serious human rights concerns. Second, there are other rights-respecting ways to protect people from harmful content. Third, extra-jurisdictionary requests to block websites should be denied. Fourth, attempts by governments to seek DNS-level blocking of websites should be public whenever possible.

A. DNS-level blocking presents serious human rights concerns

Mozilla states it does 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." We agree; DNS blocking of entire sites is a blunt instrument that should be used only in exceptional circumstances.

Filtering and blocking through DNS raises significant concerns, particularly around over-blocking and the censorship of entire web domains. As has been repeatedly noted in the reports of the UN Special Rapporteur on the freedom of opinion and expression, as well as the rulings of regional and international human rights courts, the blocking of entire web domains or tools would not be compatible with international human rights law, including the established three-part test for Article 19(3) of the International Covenant of Civil and Political Rights and

related regional human rights instruments. Indeed, this has been evident since 2011, as made clear by the reference to the analysis of "blocking or filtering technologies" in the 2011 Report of the UN Special Rapporteur.[7]

In 2020, Access Now proposed 26 recommendations on content governance that explicitly noted the dangers of internet infrastructure providers being required to take on increased content governance and censorship responsibilities.[8] In those recommendations on content governance and human rights, we recognised that infrastructure level blocking of content nearly had the same harmful impact on human rights as an internet shutdown, and specifically emphasized the following:

> Content moderation decisions by intermediaries acting at the infrastructure level (such as network and cloud security services) raise additional concerns. Their decisions, especially if internet infrastructure keeps consolidating, can result in rendering entire websites and services inaccessible. That is an extreme measure that should be carefully considered and evaluated, taking into account clear rules and principles of necessity and proportionality in a way similar to the considerations made by states when ordering the shutdown of entire sites or services.

Blocking a domain name nationwide or over a wider region should be an exceptional practice, and even in those rare instances when authorities implement them, the necessity, legality, and goals of these measures may pose more risks than benefits. Internet protocol (IP) address and DNS blocking can be excessive because services other than the infringing one may be using the same technical resources. These issues make it extremely difficult to punish the guilty without also punishing the innocent.

B. Other rights-respecting ways to protect users from illegal and harmful content

Mozilla asks if there are "more rights-protective and technically effective means of protecting users from illegal and harmful content exist beyond DNS-based blocking?" In most cases,

---

[7] Report of the Special Rapporteur on the promotion and protection of the right to freedom of opinion and expression, A/HRC/17/27, United Nations General Assembly (May 16, 2011), at ¶31, https://www2.ohchr.org/english/bodies/hrcouncil/docs/17session/A.HRC.17.27_en.pdf: "States' use of blocking or filtering technologies is frequently in violation of their obligation to guarantee the right to freedom of expression, as the criteria mentioned under chapter III are not met. Firstly, the specific conditions that justify blocking are not established in law, or are provided by law but in an overly broad and vague manner, which risks content being blocked arbitrarily and excessively. Secondly, blocking is not justified to pursue aims which are listed under article 19, paragraph 3, of the International Covenant on Civil and Political Rights, and blocking lists are generally kept secret, which makes it difficult to assess whether access to content is being restricted for a legitimate purpose. Thirdly, even where justification is provided, blocking measures constitute an unnecessary or disproportionate means to achieve the purported aim, as they are often not sufficiently targeted and render a wide range of content inaccessible beyond that which has been deemed illegal. Lastly, content is frequently blocked without the intervention of or possibility for review by a judicial or independent body."

[8] Eliška Pírková & Javier Pallero, *26 Recommendations on Content Governance*, Access Now, https://www.accessnow.org/cms/assets/uploads/2020/03/Recommendations-On-Content-Governance-digital.pdf.

authorities should target only the specific infringing information for removal and should do so only when it is also applying due safeguards in an impartial and competent judicial process. Additionally, technical means to block content should not be used to substitute from the need in many cases for law enforcement to take investigative action and prosecute alleged criminal activity; do not simply ban dissemination and while ignoring the source.

C. Non-jurisdictional requests to block websites

Mozilla asks how providers should treat requests to block entire sites outside their legal jurisdiction. Such requests should be met with skepticism. Providers should not ordinarily block or filter entire domains unless they demonstrate the rationale behind their considered view that they have been required to do so by a valid and appropriate applicable legal authority, such as Mutual Legal Assistance Treaties.

The guidance provided by General Comment 34 of the UN Human Rights Committee on Article 19 of the International Covenant on Civil and Political Rights is instructive in this broad regard. As it notes,

> Any restrictions on the operation of websites, blogs or any other internet-based, electronic or other such information dissemination system, including systems to support such communication, such as internet service providers or search engines, are only permissible to the extent that they are compatible with paragraph 3. Permissible restrictions generally should be content-specific; generic bans on the operation of certain sites and systems are not compatible with paragraph 3. It is also inconsistent with paragraph 3 to prohibit a site or an information dissemination system from publishing material solely on the basis that it may be critical of the government or the political social system espoused by the government.[9]

Mozilla's policies should follow the Human Rights Committee's principles.

D. Transparency regarding requests to block websites

Requests to block sites are sufficiently consequential that attempts to do so should be public. TRR providers should publicly indicate the jurisdictions from which they have received and accepted or rejected requests to block or filter domains, and the legal powers cited in such orders from courts and government institutions. These indicators should be included in a transparency report on content blocking/censorship published by those TRR providers. These providers should also ensure that government actors seek to make the information on such legal orders public wherever possible.

[9] Human Rights Committee General Comment No. 34, Article 19, at ¶43, United Nations International Covenant on Civil and Political Rights (Sept. 12, 2011), https://www2.ohchr.org/english/bodies/hrc/docs/GC34.pdf.

III. Building a better ecosystem

Mozilla asks how deployment of DoH will help increase trust in internet technology. Access Now works globally on technology issues and provides a Helpline service to help protect the digital security of at-risk communities, including journalists, activists, and human rights defenders. The communities we work with would benefit greatly from Mozilla's DoH system. In our previous work, we have seen repeated examples of how malicious activity and fraud in the domain name space is often used to target and harm civil society and human rights actors. For instance, it has been evident that sophisticated, relentless, and well-resourced adversaries have been attempting to surveil and hinder the work of civil society globally by using "fake domain" attacks along with other avenues of cyber threats.[10]

But Mozilla should avoid certain pitfalls in the deployment of DoH that could *reduce* trust in internet technologies. First, it is possible to have too many TRRs, and that many could be untrustworthy. Such a system would provide a false sense of security. This problem is exacerbated if Mozilla forces people to adopt DoH without any ability to opt-out. States with a record of troubling surveillance or digital attacks on activists, civil society, or journalists could seek to deploy or manage their own TRRs and pressure different browser developers to use them, thus negating any trust in the system. Allowing people to identify which TRRs they would like to use, as is the current practice, could help alleviate some of this issue.

Second, if we understand the system correctly, it hides DNS requests in HTTPS traffic. However, single-service entities or DNS servers would likely be identifiable by IP address or traffic patterns, and thus oppressive governments could identify which servers are only providing DoH service, and then use their resources and infrastructure to block DoH requests from reaching those servers or block responses back from them to the browser. Mozilla should seek to provide comprehensive "best practice" guidelines that outline how TRRs can implement the DoH service in a way that makes it more resilient against state actor abuse.

Third, assuming Mozilla continues to allow DoH opt-out, ISPs or other online actors may attempt to persuade people, against their best interests, to opt out of DoH. These actors may rely on a variety of reasons, from (unreasonable) fear that filtering services will cease functioning to the bottom-line desires of ISPs to engage in DNS-level data collection for marketing and advertising purposes. Mozilla should combat such attempts by ensuring people have information available to explain why they should decide themselves and why any outside entity (like an ISP) may be self-interested in getting people to opt-out.

---

[10] *See, e.g.*, *One of these things is not like the other*, Access Now, https://www.accessnow.org/cms/assets/uploads/archive/docs/FakeDomainsReport.pdf; *The weakest link in the chain*, Access Now (Nov. 2011), https://www.accessnow.org/cms/assets/uploads/archive/docs/Weakest_Link_in_the_Chain.pdf.

IV. Conclusion

Access Now supports Mozilla's DoH system, and submits these comments to help improve it. We look forward to consulting further.

Respectfully submitted,

Eric Null
US Policy Manager

Raman Jit Singh Chima
Asia Pacific Policy Director and Senior International Counsel

Gustaf Björksten
Chief Technologist