Response to Mozilla DNS over HTTPS (DoH) and Trusted Recursive Resolver (TRR) Comment Period

January 19, 2021

By **Gurshabad Grover** and **Divyank Katira**

The Centre for Internet and Society, India

Preliminary

This submission presents a response by the Centre for Internet & Society (CIS) to Mozilla's DNS over HTTPS (DoH) and Trusted Recursive Resolver (TRR) Comment Period[1] (hereinafter, the "Consultation") released on November 18, 2020. CIS appreciates Mozilla's consultations, and is grateful for the opportunity to put forth its views and comments.

Responses to issues for consultation

Respecting security and privacy

Q1. 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. To what extent can our requirement be shortened further while allowing providers sufficient data to operate the service? What operational constraints, if any, are created by this maximum 24-hour retention time?

We are not aware of any operational constraints that are created by the maximum retention time of 24 hours.

Mozilla can also explore ways of nudging users to pick DNS providers that keep data ephemerally, or at the very least signal this information to the users. For instance, DNS providers that keep data ephemerally can be identified in Firefox's interface with an appropriate visual indication for offering more privacy than other resolvers in the list.

Q4. 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?

In addition to the insight about governmental requests, we recommend expanding the transparency requirements to content filtering orders and practices.

Generally, the requirement of maintaining transparency reports can be significantly improved by providing guidance to resolvers on best practices on the same that have been recommended by civil society. For instance, the policy can note that the transparency reports must mandatorily include:[2]

---

[1] Owen Bennett and Udbhav Tiwari, "Mozilla DNS over HTTPS (DoH) and Trusted Recursive Resolver (TRR) Comment Period: Help us enhance security and privacy online", Mozilla Open Policy & Advocacy Blog, 18 November 2020, <https://blog.mozilla.org/netpolicy/2020/11/18/doh-comment-period-2020/>

[2] *We have incorporated advice from three civil society efforts that recommend best practices for transparency reports:* Nate Cardozo, et al, "Who Has Your Back? Censorship Edition 2018", Electronic Frontier Foundation, <https://www.eff.org/who-has-your-back-2018>; Multiple authors, "The Santa Clara

- A public document that clearly specifies that the resolver only acts on reasoned legal orders from state authorities, and restricts the geographical scope of the blocking as much as possible[3]
- The number of blocking requests received by the resolver by governments and third parties (and optionally statistics on reasons given in the written requests)
- The number of information requests received by the resolver from state authorities (and optionally statistics on reasons given in the written requests)
- The blocking decisions and content filtering policies made by the resolver out of its own volition

Online safety

Q1. 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?

Considering the growing use of third-party DNS resolvers, it is entirely possible that a resolver receives a blocking request from a jurisdiction that it is not based in. For the continued availability of the service in that jurisdiction, the resolver may have to comply with that request. We believe that such compliance should not disqualify a resolver from being considered as a TRR.

Rather, a proportionate approach may be to require all resolvers to restrict the geographical scope of the blocking as much as possible to the jurisdiction from which it received the blocking request. This will avoid another pitfall of the current framing: it would ensure that users of the resolver that are not based in the jurisdiction (that the resolver is based in) continue to receive accurate DNS responses.

Additionally, Mozilla should disqualify resolvers that do not make their best efforts to limit their blocking to the geographical scope of the jurisdiction from where it received the blocking request.

Q2. What harmful outcomes can arise from filtering/blocking through the DNS?

If done without proper disclosures (like block notices), blocking through DNS can be misdiagnosed as an outage or malicious attack, and contribute to unnecessary

---

Principles On Transparency and Accountability in Content Moderation" Santa Clara Principles (2018), <https://santaclaraprinciples.org/>; Torsha Sarkar, Suhan S and Gurshabad Grover, "Through the looking glass: Analysing transparency reports", Centre for Internet and Society (31 October 2019), <https://cis-india.org/internet-governance/files/A%20collation%20and%20analysis%20of%20government%20requests%20for%20user%20data%20%20and%20content%20removal%20from%20non-Indian%20intermediaries%20.pdf>

[3] *For example, if a resolver receives a request from the Indian Government to block example.com, it should still serve the accurate DNS response to requests received from 'non-Indian' IP addresses.*

troubleshooting.[4] By itself, DoH also does not provide an important security property: authentication.[5] Thus, we believe that any filtering through DNS relies on giving an *inauthentic* response to users, and should be seen as exploiting a security vulnerability in DNS protocols. Apart from the general impediment the use of such blocking causes to the deployment and use of protocols like DNSSEC, it contributes to an erosion of public trust in internet infrastructure.[6]

Q3. What more rights-protective and technically effective means of protecting users from illegal and harmful content exist beyond DNS-based blocking?

Governments, internet service providers and intermediaries have a number of ways to regulate illegal and/or harmful content. Rather than relying on intermediaries, we should move towards endpoint-based blocking, i.e. filtering decisions made at user devices.[7] Instead of imposing content decisions, such an approach empowers users. Compared to blocking decisions that happen through other nodes in the network, "endpoint-based blocking is the least likely to cause collateral damage to Internet services or the overall Internet architecture."[8] This is because end devices can see into all layers and content of the communication and thus define the narrowest-possible filters. The unintended consequences of such blocking are also minimised to the endpoint. While a complete shift towards end-point based blocking is a broad and long-term mission, organisations and individuals can disincentivize network-based blocking because of the disproportionate nature of the harms it causes.

Q4. 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.) 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? What challenges weigh against a requirement to publish block lists?

While we support the TRR policy's direction for trusted resolvers to publish block lists, such a requirement can be at odds with laws in different jurisdictions. For instance, one of the laws that govern the procedure for the Government of India to issue blocking requests to online

---

[4] Security and Stability Advisory Committee, "SSAC Advisory on Impacts of Content Blocking via the Domain Name System", October 2012, ICANN, <http://www.icann.org/en/groups/ssac/documents/sac-056-en.pdf>

[5] *Unless used with DNSSEC.*

[6] Security and Stability Advisory Committee, "SSAC Advisory on Impacts of Content Blocking via the Domain Name System", ICANN (October 2012), <http://www.icann.org/en/groups/ssac/documents/sac-056-en.pdf>; R. Barnes, et al, "RFC 7754: Technical Considerations for Internet Service Blocking and Filtering", Internet Architecture Board (March 2016), <https://tools.ietf.org/html/rfc7754>; Internet Society, "Internet Society Perspectives on Internet Content Blocking: An Overview", Internet Society (March 2017), <https://www.internetsociety.org/resources/doc/2017/internet-content-blocking/>

[7] Id (R. Barnes, et al)

[8] Id

intermediaries obligates intermediaries to keep such requests confidential.[9] Thus, the requirement of publishing blocklists may preclude resolvers operating in certain jurisdictions from joining Mozilla's TRR program. To meet the objectives of both transparency and geographical diversity, the audit requirements can include a legal justification for why complete transparency of blocking practices is not possible.

Q5. How can we best present information about opt-in filtering endpoints to end users (e.g., for malware blocking or family-friendly blocking)?

In the list of available resolvers, Firefox can visually indicate filtering practices associated with each resolve. These can include:

- whether a resolver filters domains based on state requests
- whether a resolver filters domains based on third-party requests, including copyright
- whether a resolver filters domains based on their own policies (malware, family-friendly blocking, etc.)

These can be accompanied with a longer explanation of the filtering policies.

Misc.: Responses for blocked domains

While not a question for consultation, we noted that the policy recommends that the resolver respond with NXDOMAIN when a domain requested by a user is not present. Since this requirement is in the same section as 'blocking', we recommend that the policy clarify that this is *not* applicable to circumstances when a resolver is actively filtering content.

If and when a resolver is blocking a domain name, a NXDOMAIN response can be cause for confusion for users.[10] We recommend that domain filtering be done in a transparent way: the Mozilla TRR policy can recommend resolvers to respond to a 'blocked' domain with an IP address that leads to a block notice, i.e. a page that identifies the organisation operating the resolver, and the reason for why the domain is being blocked.

Building a better ecosystem

Q2. What exploitations of the DNS in your region could DoH protect against?

[9] Rule 16, Information Technology (Procedure and Safeguards for Blocking for Access of Information by Public) Rules, 2009. <https://cis-india.org/internet-governance/resources/information-technology-procedure-and-safeguards-for-blocking-for-access-of-information-by-public-rules-2009>

[10] *See* Kushagra Singh, Gurshabad Grover and Varun Bansal, 'How India Censors the Web', Proceedings of the 12th ACM Conference on Web Science, <https://dl.acm.org/doi/abs/10.1145/3394231.3397891> *which found that Bharti Airtel is censoring web content by responding to certain DNS queries with NXDOMAIN, and notes this as contributing to opaqueness in blocking.*

In India, DNS injection and poisoning are both employed by local ISPs to enforce Government censorship orders.[11] The deployment of DoH could potentially act as a tool for circumvention of such censorship. There is also evidence of unlawful and arbitrary blocking by ISPs in India.[12]

Apart from minimal security requirements, there is no modern data protection law in operation in India. Telecom and internet service providers in India have broad privacy policies with very little safeguards,[13] and some of them may be profiteering from data on user behaviour,[14] which may potentially include DNS queries.

All these harms can be mitigated by large-scale deployment of encrypted DNS protocols, provided that the choice of DNS resolver does not default to the ISP's.

Q3. What are the best ways to gain global adoption/support of the DoH standard amongst ISPs and DNS providers?

The proposals being discussed at the Adaptive DNS Discovery (ADD) working group at the Internet Engineering Task Force (IETF) have a clear direction: how to discover resolvers, *within the network*, that supported encrypted queries.[15] The standardisation of these proposals will go a long way in encouraging ISPs to deploy their own DoH/DoT resolvers.

It is appreciable that Mozilla is working with ISPs like Comcast to devise ways for local encrypted DNS resolver discovery,[16] but we should also note that ISPs in other jurisdictions (including India) are more likely than other resolvers to have censorship practices. In working with other stakeholders to encourage DoH deployment, we urge Mozilla to not fall into the trap of de-prioritising user choice of DNS.

General comments regarding TRR policies

[11] Id.
[12] Id.
[13] Internet Freedom Foundation, "Privacy Policies of Telecom Service Providers - Or Why You Shouldn't Just Click Accept", Internet Freedom Foundation (December 2020), <https://internetfreedom.in/privacy-policies-of-telecom-service-providers-or-why-you-shouldnt-just-click-accept/>
[14] Promit Mukherjee, "From big oil to big data: inside Mukesh Ambani's $20 billion start-up", Reuters (August 2016), <https://www.reuters.com/article/reliance-telecoms-jio-idINKCN11611V?edition-redirect=in>
[15] Gurshabad Grover, "Adapting to the reality of encrypted DNS deployment", Council of European National Top-Level Domain Registries (December 2020), <https://centr.org/news/news/ietf109-encrypted-dns.html>
[16] Eric Rescorla and Jason Livingood, "CNAME Discovery of Local DoH Resolvers", *Work in progress Internet Draft*, <https://datatracker.ietf.org/doc/draft-rescorla-doh-cdisco/>; Eric Rescorla, "More details on Comcast as a Trusted Recursive Resolver", Mozilla Blog (June 2020), <https://blog.mozilla.org/blog/2020/06/26/more-details-on-comcast-as-a-trusted-recursive-resolver/>

- **More clarity on purpose of user data collection:** The TRR policy allows, under certain conditions, for the collection of both aggregate and identifiable user data for “the purpose of operating the service”. Although it appears that the intent of this statement is to limit collection to operational metrics that are required to efficiently run a resolver, such a broad framing can be construed in a number of different ways. It may be beneficial to define this more narrowly. For instance, the purpose of user data collection can be constrained to “tuning performance of the resolver and debugging technical issues”.

- **Privacy notice should disclose how aggregate data is used:** The transparency requirements outlined in the policy require a trusted resolver to document specific fields for aggregate data that will be retained but not how this aggregate data will be used. Given the large amounts of sensitive user information a trusted resolver will encounter, the data may be valuable and offer a competitive advantage even in a de-identified, aggregate form. Individuals must be informed of all potential uses of such data collected from them.

- **Transparency around excluding resolvers from the TRR program:** Since the decision to include and exclude parties from the TRR program is at Mozilla’s sole discretion, in addition to violations of the policy, Mozilla should also publicly document instances where resolvers were not allowed to join the TRR program accompanied by a relevant explanation.