New America's Open Technology Institute Submission to Mozilla Comment Period on DNS-over-HTTPS Implementation

January 2021
Ross Schulman
Spandana Singh

Respecting Privacy and Security

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).

The Policy should permit exceptions to allow TRRs to provide data to law enforcement in situations where their data requests meet the required legal definition of an emergency circumstance.

There may be some other limited situations in which a TRR should be allowed to collect additional data or retain data for longer than the 24 hours stated in the resolver agreement. Narrow exceptions should also be permitted to enable TRRs to secure either the resolver's own infrastructure or the security of the internet more broadly. Specifically, resolvers should be permitted to retain or collect records when:

- The resolver's own infrastructure is under cyber attack and the data collected or retained may be necessary for mitigating, investigating, or remediating the attack and its aftermath.
- The resolver has a good faith belief that DNS queries they are receiving in real time reflects an ongoing widespread attack, and they need to retain data for longer than the standard 24 hour period for the purposes of mitigating the attack. This would include sharing data with other infrastructure providers or a Computer Emergency Readiness Team (CERT) as needed to respond to the attack, but limited to only those records which are necessary to understand and address the attack. The policy should require that data be de-identified to the extent possible before any such data sharing.

In these situations, resolvers should publish a publicly available transparency report which, at a minimum, outlines:

- A clear explanation of the emergency circumstances in which the exemption applies to permit the resolver to provide data. This should include explanation of emergency circumstances in which a resolver is required by law to provide data to the government and other emergency circumstances including when a resolver has identified an ongoing attack

- The number of government requests for data under the emergency exception to the TRR policies the resolver has received
- The percentage of these emergency circumstance requests that were accepted by the resolver and the percentage that were rejected by the resolver
- A break down of the total number of emergency circumstance requests the resolver has received based on the category or type of emergency circumstance
- A break down of the total number of emergency circumstance requests the resolver has received by country
- The total number of times a resolver has engaged in additional data collection due to an ongoing and widespread attack
- A break down of the total number of times a resolver has engaged in additional data collection due to an ongoing and widespread attack based on the category or type of attack

In addition, Mozilla itself should also publish a publicly available transparency report which at a minimum outlines:

- Which resolvers are participating in the TRR program
- The total number of Mozilla Firefox users that have turned DNS-over-HTTPS on or off
- A break down of the total number of Mozilla Firefox users that have turned DNS-over-HTTPS on or off by country
- The total number of enterprise or network operators that have turned off DNS-over HTTPS
- A break down of the total number of enterprise or network operators that have turned off DNS-over-HTTPS by country

As OTI has outlined in our Transparency Reporting Toolkit on U.S. Government Requests for User Information and our Transparency Reporting Toolkit on Content Takedowns, reporting entities should issue transparency reports on clearly and consistently delineated reporting periods and in an openly licensed, machine-readable format. In instances where the reporting entity is acquired or reorganizes its website, it must ensure that the transparency report maintains a functioning URL. Transparency reports for different reporting periods should also be housed in a central location online to enable easy access and comparison.

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?

As OTI has outlined in our Transparency Reporting Toolkit on U.S. Government Requests for User Information, entities that receive government requests for data should publish publicly available transparency reports in order to provide adequate transparency and accountability around the requests for data that they receive. As noted above, these reports should include

data regarding occasions on which the resolver collects additional information under emergency circumstances, including government requests for data on an emergency basis and situations in which the resolver is responding to an ongoing and widespread cyber attack. In addition, at a minimum, resolvers should publish a transparency report which outlines:

- **Explanation of legal processes:** Resolvers should provide clear and comprehensive explanations of legal processes that underpin any government requests for data they receive. Definitions or a glossary explaining legal processes and other key terms used in the report can inform readers about the types of processes that might allow governments to access their data, educate readers on the different legal processes the resolver requires to turn over specific types of information, and generate an understanding of the logistics behind transparency report, including how companies count legal processes.

- **Reporting on scope and scale of different legal processes:** Resolvers should provide clear and granular categorizations of applicable legal processes that underpin any government requests for data they receive. In addition, an ideal report will, at minimum, provide the number of government requests received under each of these categories.

- **Reporting on the subjects of requests and how users are impacted:** Resolvers should report on the number of selectors specified in a request. Where relevant, the resolver should also report the number of users and/or accounts implicated in a request, making a clear distinction when both are implicated.

- **Reporting on outcomes and compliance with requests:** Resolvers should report on the percentage of government requests that they accepted and disclosed information under, the percentage of government requests they rejected, etc. These data points should be broken down according to the relevant legal process.

- **Reporting on user notification:** When the TRR has another relationship with the user (e.g. it serves as that user's Internet Service Provider) that could allow them to connect individuals with their DNS queries, resolvers should provide clear, comprehensive, and granular reporting on notification to users. This includes reporting on three types of notifications: 1) When a request was under seal and the user could be notified, 2) When a request was not under seal and the user was notified, and 3) When a request was not under seal and the user was not notified.

Resolvers should issue these transparency reports on clearly and consistently delineated reporting periods and in an openly licensed, machine-readable format. In instances where the reporting entity is acquired or reorganizes its website, it must ensure that the transparency report maintains a functioning URL. Transparency reports for different reporting periods should also be housed in a central location online to enable easy access and comparison. Further, transparency reports should reflect the unique circumstances and offerings of resolvers. However, we encourage a degree of standardization among resolver reports, as this enables

third-parties such as researchers, civil society, and journalists, to compare and analyze the data in these transparency reports.

In addition to publishing transparency reports, resolvers should establish clear policies for processing and responding to government requests for data. These policies should include guidelines on the following:

- **Tracking requests and their status:** Resolvers should use a single, centralized process for tracking, tagging and keeping tabs on the status of government requests from the moment they are received until the time a response is provided to the government.

- **Reviewing and classifying requests for accuracy and validity:** Before a resolver can respond to a request, it has to identify the type of process and the agency or court that issued it. Resolvers should therefore have trained staff that can properly classify government and legal requests. This process is also vital for identifying requests that do not comply with legal requirements and requests that companies should push back on.

- **Responding to requests:** Resolvers should work with their legal counsel to develop a playbook for responding to government requests before they receive them. This will prevent errors and promote greater safeguarding of privacy and security.

- **Providing user notice:** Notifying users when their information has been requested is an important aspect of safeguarding user rights. However, at times the provision of notice can be delayed for ongoing investigations subject to applicable laws. In some circumstances the law permits the government to issue a gag order that prevents companies receiving data requests from informing the target of the investigation. For those requests without gag orders, when the TRR has another relationship with the user (e.g. it serves as that user's Internet Service Provider) that could allow them to connect individuals with their DNS queries, resolvers must decide whether, how, and under what circumstances they'll provide notice to their users. For requests with gag orders, resolvers must decide whether to challenge the order and/or to inform users after the gag order has been lifted.

- **Keeping data secure:** Information about law enforcement and intelligence requests is sensitive information in itself. In order to keep this information secure, resolvers should carefully consider how this data is maintained and who has access to it.

- **Challenging requests where necessary and appropriate:** Resolvers should have procedures in place that enable them to evaluate the validity and accuracy of requests and challenge requests that are out of scope or that infringe on privacy and security.

A public version of these policies should be posted online for users and other parties to view. This public edition of the policies should include information on what kind of legal orders and

mechanisms the resolver accepts and responds to, the format in which requests must be made, the scope requests must fall under and, as applicable, the user notification process. Providing a FAQ section which augments understanding of how companies respond to government requests for data is also helpful.

Online Safety

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

Many policymakers and advocates are placing increased pressure on tech companies to address the spread of illegal and harmful content online. As a result, some companies have adopted DNS filtering and blocking practices. OTI recognizes the importance of addressing illegal and harmful content online. However, except in the limited situation described below regarding malware command and control domains, DNS filtering and blocking are not effective methods of tackling these forms of content. They also raise numerous concerns related to infringements on freedom of expression rights and on users' right to access, seek, and exchange information and ideas.

DNS blocking and filtering particularly raises freedom of expression concerns as it results in an entire domain name being blocked, rather than a specific page within a website. For example, if copyrighted content was identified on a website, and DNS blocking and filtering was implemented, the entire website would be blocked, rather than just the infringing content. As a result, a significant amount of non-infringing speech would be impacted. The impact of this practice on freedom of expression has been recognized even in the context of efforts to address the harms of child sexual abuse materials. For example, in February 2011, the U.S. Immigrations and Customs Enforcement seized the mooo.com domain name for owning or distributing child pornography. By blocking the mooo.com domain name, ICE authorities erroneously shut down 84,0000 websites for several days, and falsely implicated them as committing a serious criminal offense. Eventually, authorities reversed this action, but this case study demonstrates the far-reaching impacts DNS blocking and filtering can have on speech and content online.

This recognition that DNS blocking is an overbroad tool was reinforced in Center for Democracy & Technology v. Pappert, in which the court struck down the Pennsylvania Internet Child Pornography Act (ICPA), which incorporated DNS filtering. The ICPA would have allowed the Pennsylvania Attorney General or any District Attorney in the state to seek a court order based on probable cause that required an internet service provider (ISP) to disable access to child pornography within five days or face criminal liability. However, the court found that the ICPA was not in line with the First Amendment, as the use of DNS filtering practices would result in the over-blocking of content that did not violate child pornography laws, and the harms caused by this overbroad censorship outweighed the legitimate need to address child pornography online.

DNS filtering and blocking are also not effective methods for addressing harmful and illegal

content online as the individuals responsible for uploading the content may be based in a different country than the entity carrying out the blocking. As a result, the two parties would be subject to different laws, which may differently define "illegal content". In addition, DNS filtering and blocking is only effective when the blocking entity retains control over the entire network connection of the end user. If a user has the option to use a different connection or a different set of DNS servers, they can avoid the content blocking or filtering altogether. For example, in 2012, Turkey blocked several DNS queries, but users were able to avoid the block by changing their systems to use Google's public DNS servers.

As noted, there is one exception where resolver blocking of domain names may have a net-benefit on the overall security of the internet without meaningfully impacting users' freedom to access the whole of the internet: malware command and control domains. After infecting a user's computer, malware often contacts a server operated by its own authors for further instructions or to exfiltrate stolen data. While end-point security options can serve to block those connections, blocking the DNS requests serves to protect everyone.

There are two reasons why blocking of these kinds of DNS queries poses less of a problem than other DNS blocking: First, these domains are usually random strings of letters and numbers that don't serve any actual content. Blocking them will not prevent legitimate users from visiting any actual content that they might want to see. Secondly, the universe of these types of domains is finite and definable. Unlike trying to identify "adult" or "blasphemous" content, it should be possible to rely on a trusted third party to maintain a list of domains serving as command and control servers at any given time that TRRs could use as a blocklist, and to encode that list into specific contractual language.

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

As discussed above, DNS filtering and blocking practices are methods of addressing illegal and harmful content online that are ineffective and that raise numerous threats to freedom of expression online. We recommend instead that entities that host user-generated content engage in rights-respecting content moderation operations. We note that contemporary content moderation practices are in no way perfect. We have continuously expressed concerns related to the increasing use of unreliable automated tools and unclear content moderation policies and have pushed companies to improve these efforts. However, in comparison to the blunt tools available to resolvers, relying instead on platforms that host user-generated content or end-user device moderation tools to adopt responsible content moderation policies and practices is a better solution to address the spread of harmful and illegal content online at scale without causing as much collateral damage to online speech.

In addition, although the content moderation policies and practices of internet companies are problematic in many ways, companies' moderation policies and practices are far more transparent than DNS filtering and blocking policies and practices. Today, there is very little transparency and accountability around how and when DNS filtering and blocking practices are

used, and what company policies guide their adoption. The opaque nature of this practice raises risks that it can be easily abused to censor content, therefore further threatening freedom of expression.

In order to ensure that content moderation efforts are right-respecting, companies that host user-generated content or end-user device moderation tools must consult with a broad range of stakeholders from around the world, including representatives from human rights groups, civil rights organizations, and civil society when crafting their policies and practices. In addition, companies must provide adequate transparency and accountability around their content moderation operations. We have provided guidance on this in the Santa Clara Principles on Transparency and Accountability in Content Moderation, our Transparency Reporting Toolkit on Content Takedowns, and our report *Everything in Moderation: An Analysis of How Internet Platforms Are Using Artificial Intelligence to Moderate User-Generated Content*.

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.)

In situations where TRRs must engage in legally required blocking practices, they should publish publicly available transparency reports that at a minimum outline:

- A list of the domain names that are blocked, broken down by country.
- A clear explanation of the policies that a company follows when it is legally required to block certain domain names. This should include an explanation of how a company determines whether a legal request to block a domain is sufficient and valid, and in what instances it will challenge such a request.
- An outline of in what instances the company will notify impacted users that a domain is blocked, and how this notice will be disseminated.

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?

In instances where a company creates and maintains a block list, it should offer domain owners access to a complaint and redress process. If a domain owner has their own domain placed on a block list, they should receive notification of this decision. In addition, domain owners should be able to appeal the inclusion of their domain on a block list.