Federal Register Notice 87 FR 35250, https://www.federalregister.gov/documents/2022/06/09/2022-12432/request-for-information-on-advancing-privacy-enhancing-technologies, June 9, 2022

Request for Information (RFI) on Advancing Privacy Enhancing Technologies

Cloudflare

DISCLAIMER: Please note that the RFI public responses received and posted do not represent the views or opinions of the U.S. Government. We bear no responsibility for the accuracy, legality, or content of the responses and external links included in this document.

VIA EMAIL
Dr. Alondra Nelson, Director
Office of Science and Technology Policy
Eisenhower Executive Office Building
725 17th Street NW
Washington, D.C.

Re: RFI response: Privacy-Enhancing Technologies, FR Doc. 2022-12432

Dear Dr. Nelson,

We write in response to the Request for Information (RFI) published by the Office of Science and Technology Policy seeking comment on Privacy-Enhancing Technologies (PETs). Cloudflare appreciates the opportunity to comment, and we hope the exercise will result in a useful sharing of information amongst stakeholders.

Cloudflare has faced a variety of challenges in seeking to deploy PETs, ranging from industry resistance to restricting data collection to differing views among governments about new technology. We believe the U.S. government can play an important role in helping both to encourage adoption of new PETs and to build consensus between governments about how to address new PETs.

Cloudflare submits the following comments, which will address our own experience in PETs, including collecting data in privacy-enhancing ways; a brief discussion of research opportunities, risks, and limitations of PETs; our views on the regulation of PETs; and our recommendations for further work to advance and support the development of PETs.

Introduction and Cloudflare Background

Cloudflare is an Internet performance and security company that is on a mission to help build a better Internet. Privacy-enhancing technologies are critical components to achieve that goal. We operate a global network of points of presence in 270+ cities, within 50ms of 95% of the world's population. We use that network to provide services that help people increase the security and performance of their web sites and services.

Cloudflare's Work on PETs

Since its inception in 2010, Cloudflare has invested in the development and deployment of cutting-edge cryptography and Internet standards to help people improve the security and performance of their websites and services. Our company's goal is to enable a research

environment that generates both new knowledge and technology that can lead to innovative products. Cloudflare's research team[1] works hand-in-hand with both product and engineering to help drive long-term positive outcomes for both Cloudflare's customers and the Internet at large. Our researchers are active contributors to the Internet Engineering Task Force[2] (IETF) and other standards organizations, and our extensive network allows us to bring privacy-enhancing technologies from the experimental stage through deployment in our extensive network.

We've released a number of services that deploy state-of-the-art, privacy-enhancing technologies for DNS and other communications to help individuals, businesses, and governments alike, and we've made substantial contributions to the development of privacy-enhancing Internet protocols. Here are a few notable examples of our work in the area of PETs:

- **DNS Privacy:** Since 2018, Cloudflare has maintained a free DNS resolver (1.1.1.1) that powers approximately 16.1 million DNS queries per second. 1.1.1.1 supports private and encrypted requests from clients via DNS over HTTPs (DoH) and DNS over TLS (DoT), providing confidentiality for DNS traffic in transit to Cloudflare. DNS requests can contain some sensitive data, such as your location, the domains and subdomains you have visited, the time of day requests were submitted, and how long you stayed on certain sites. The 1.1.1.1 resolver never writes the querying IP addresses together with the DNS query to disk, and unless someone is using 1.1.1.1 for Families, we don't receive the EDNS Client Subnet Header (which might reveal location information).

  1.1.1.1 also supports other privacy-enhancing mechanisms for DNS, including query name (QNAME) minimization and omission of EDNS Client Subnet extensions in DNS messages, both of which limit information about what and who is requesting resolution of a particular name to authoritative servers.

  Even with these standard mechanisms deployed, 1.1.1.1 is operated by a single entity that can still technically link all queries to client IP addresses. As a result, Cloudflare went beyond the basics to develop and deploy Oblivious DNS over HTTPS (ODoH),[3] which separates IP addresses from queries, so that no single entity can see both at the same time.[4]

- **Transport Layer Security (TLS)/QUIC**: Cloudflare has played an active role in developing transport security protocols in wide use on the Internet, from our initial deployment of TLS 1.3 in 2014,[5] to 2018 when we made it the default setting on all of

[1] For background on Cloudflare's research team, including research areas, staff profiles, and publications, see https://research.cloudflare.com/.
[2] See The Internet Engineering Task Force (hereinafter, IETF) at https://www.ietf.org/.
[3] See IETF, Oblivious DNS over HTTPS, at https://datatracker.ietf.org/doc/rfc9230/.
[4] See E. Kinnear, P. McManus, T. Pauly, T. Verma & C.A. Wood, Oblivious DNS over HTTPS, IETF Independent Submission, June 2022, available at https://www.rfc-editor.org/rfc/rfc9230.pdf.
[5] See Matthew Prince, Introducing Universal SSL, The Cloudflare Blog, Sept. 29, 2014, at https://blog.cloudflare.com/introducing-universal-ssl/.

our customer sites.[6] Cloudflare was also one of the early adopters of QUIC and is today one of the largest deployments.

- **Encrypted Client Hello (ECH)**: ECH is a new extension for TLS that will significantly improve its privacy. Today, a number of privacy-sensitive parameters of the TLS connection are negotiated in the clear. This leaves a trove of metadata available to network observers, including the endpoints' identities, how they use the connection, and so on. ECH encrypts the full handshake so that this metadata is kept secret. Together with major web browsers and other stakeholders, Cloudflare has helped develop the standard and support initial experiments.

- **Privacy Pass**: Cloudflare developed and released the privacy-enhancing technology called Privacy Pass in 2017,[7] and it has since evolved into a standard used by many website operators (it's on track[8] to be adopted as a standard by the IETF). Privacy Pass lets users prove their identity across multiple sites anonymously without enabling tracking. When people use anonymity services or shared IPs, it makes it more difficult for website protection services like Cloudflare to identify their requests as coming from legitimate users and not bots. To help reduce the friction for these users — which include some of the most vulnerable users online — Privacy Pass provides them with a way to prove they are legitimate across multiple sites on the Cloudflare network. This is done without revealing their identity, and without exposing Cloudflare customers to additional threats from malicious bots.

- **Distributed Aggregation Protocol (DAP)[9]:** Cloudflare researchers have been instrumental in developing DAP, a new protocol to enable privacy-preserving measurement. The protocol involves a large set of clients and a small set of servers and is used to compute aggregate statistics over the clients' inputs without learning the inputs themselves. DAP enables the use of recent advances in multi-party computation in order to ensure that no input is ever seen in the clear by any server. Several organizations are actively working on DAP. For example, the Internet Security Research Group[10] (ISRG) recently raised one million dollars to fund Divvi Up, a system for privacy-respecting aggregate statistics.[11] Apple and Google released a white paper last

[6] See Alessandro Ghedini, You get TLS 1.3! You get TLS 1.3! Everyone gets TLS 1.3!, The Cloudflare Blog, May 16, 2018, at https://blog.cloudflare.com/you-get-tls-1-3-you-get-tls-1-3-everyone-gets-tls-1-3/.
[7] See Nick Sullivan, Cloudflare Supports Privacy Pass, The Cloudflare Blog, Nov. 9, 2017, at https://blog.cloudflare.com/cloudflare-supports-privacy-pass/.
[8] See Tommy Pauly, Steven Valdez, & Christopher A. Wood, The Privacy Pass HTTP Authentication Scheme, IETF, July 6, 2022, at https://datatracker.ietf.org/doc/draft-ietf-privacypass-auth-scheme/.
[9] See Tim Geoghegan, Christopher Patton, Eric Rescorla, & Christopher A. Wood, Distributed Aggregation Protocol for Privacy Preserving Measurement, IETF, May 4, 2022, at https://datatracker.ietf.org/doc/draft-ietf-ppm-dap/.
[10] See Internet Society Research Group, at https://www.abetterinternet.org/.
[11] See Dan Fernelius, ISRG Raises More Than $1M For Advancing Divvi Up, Divvi Up, June 28, 2022, at https://divviup.org/blog/2022-06-28-announcing-divviup-funding/.

year explaining how their COVID-19 exposure notifications system was based on a predecessor of the new DAP standard.[12]

Deep Dive: Standardizing Multi-Party Computation for Privacy-Preserving Measurement (RFI topic #1)

A type of privacy-enhancing technology that OSTP identified as an area of particular interest is secure multi-party computation. Within that field, privacy-preserving measurement (PPM) is a critical operation that refers to the measurement of digital data in ways that are privacy-preserving for common tasks such as sharing, collaboration, and analysis. In its consideration of specific research opportunities to advance PETs (topic #1 of the RFI), PPM should be one of the areas that OSTI includes as part of its national strategy. Organizations and researchers often want to measure certain operations in which sensitive or personal data are present. For example, an Internet infrastructure company such as Cloudflare might want to measure how long it takes for clients to render a web page served by Cloudflare. Or a medical research institution or public health agency might want to measure the spread of infection in a community. In both of these situations, the measuring entity's interest lies in aggregated data, not in individual data points or identities themselves. By using PPM techniques, the measuring entity can calculate the aggregated data and associated values in a way that avoids any exposure of the personal or sensitive data.

In early 2022, IETF established the PPM working group.[13] The objective of the working group is to investigate and develop protocol standards for aggregating user measurements in privacy-preserving ways, including the deployment of new cryptographic techniques. The goal is to ensure that the server (or multiple, non-colluding servers) can compute the aggregated value without learning the value of individual measurements. To do so, the IETF PPM working group will standardize protocols for deployment of these techniques on the Internet.

Cloudflare researchers played an integral role in forming the IETF PPM working group. Cloudflare continues to help guide the standardization process by co-editing documents, guiding the agenda in meetings, and developing or helping to maintain open-source implementations.

The direction of the work of the IETF's PPM working group is determined primarily by use cases brought to the working group. So far, this includes applications like web browser telemetry, web analytics, telemetry for COVID-19 exposure notification platforms,[14] and various machine learning tasks. Any organization, be it a private company or a public institution, is free to bring use cases to the group and participate in the protocol design process.

---

[12] See Apple & Google, Exposure Notification Privacy-Preserving Analytics (ENPA) White Paper, Apr. 2021, available at https://covid19-static.cdn-apple.com/applications/covid19/current/static/contact-tracing/pdf/ENPA_White_Paper.pdf.

[13] See Privacy Preserving Measurement (PPM), IETF, at https://datatracker.ietf.org/wg/ppm/about/.

[14] See Apple & Google, supra note 12.

Like many IETF working groups, the IETF PPM working group is made possible by engagement with the academic community. Through journal articles, conference papers, and other artifacts, researchers provide critical insights into the (in)security of proposed protocols and, especially in the case of PPM, devise solutions for specific problems. We have seen a number of papers in recent years proposing multi-party computation (MPC) schemes for specific (classes of) secure data aggregation tasks. These works draw from and build on a wide variety of sub-disciplines of cryptography, including zero-knowledge proof systems, (partially) homomorphic encryption, oblivious transfer, and others.

In many of these systems,[15] the computation of the aggregate is distributed amongst multiple servers as follows: Clients "split" their measurements into multiple "secret shares" and send one share to each server. Once all the shares have been uploaded, the servers use them to reconstruct the desired aggregate result. Splitting the measurements into secret shares ensures they are kept secret: As long as at least one server executes the protocol honestly, no individual's measurement is ever observed in the clear. At the same time, many of these systems are designed to ensure that the result is always computed correctly, even in the presence of malicious (or merely misconfigured) clients that would otherwise provide invalid inputs to the computation.

This paradigm improves significantly upon the status quo. While a number of solutions exist for improving privacy of the measurement aggregation process, none other than MPC allows the aggregate result to be computed both precisely and without revealing measurements to the data collector.[16] MPC also provides an opportunity for a third-party audit of the collection process, thereby ensuring much stronger security performance.

The "cost" of stronger security is additional operational complexity, since the reliability of the data collection system now depends on multiple servers instead of just one. Addressing this complexity is one of the principal tasks of the IETF PPM working group. The DAP standard mentioned above is the first draft taken up by the IETF PPM working group.[17] The objective of the DAP specification is to spell out precisely how to execute MPC protocols for secure aggregation over HTTPS. The specification targets a particular class of constructions, called

---

[15] See, e.g., Ulfar Erlingsson, Vasyl Pihur, & Alesandra Korolova, RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response, Aug. 25, 2014, available at https://arxiv.org/pdf/1407.6981.pdf; Henry Corrigan-Gibbs & Dan Boneh, Prio: Private, Robust, and Scalable Computation of Aggregate Statistics, Proceedings of the 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI '17), Mar. 27-29, 2017, Boston, MA, available at https://www.usenix.org/system/files/conference/nsdi17/nsdi17-corrigan-gibbs.pdf; Dan Boneh, Elette Hoyle, et al., Lightweight Techniques for Private Heavy Hitters, 2021 IEEE Symposium on Security and Privacy (SP), 2021, pp. 762-776, available at https://eprint.iacr.org/2021/017.pdf; and Joseph J. Pfeiffer, Denis Charles, Davis Gilton, et al., Masked LARk: Masked Learning, Aggregation and Reporting worKflow, Oct. 27, 2021, available at https://arxiv.org/pdf/2110.14794.pdf.

[16] Differentially private systems such as RAPPOR [EPK14] in which clients randomize their uploaded measurements, provide some degree of privacy protection, but the aggregate result is not computed precisely.

[17] See Tim Geoghegan, Christopher Patton, Eric Rescorla, & Christopher A. Wood, Distributed Aggregation Protocol for Privacy Preserving Measurement, IETF, May 5, 2022, at https://datatracker.ietf.org/doc/draft-ietf-ppm-dap/.

Verifiable Distributed Aggregation Functions (VDAFs).[18] A specification of VDAFs, including a few concrete instantiations from the literature, is under development in the CFRG[19] research group of the IRTF.[20] The goal of this work is to provide cryptography researchers with an explicit set of design criteria for developing new constructions. In turn, the DAP specification provides a clear roadmap to deployment of these constructions.

Our hope is that the work developed in the IETF PPM working group will enable organizations of all sizes to adopt a private-by-default posture for their data collection and aggregation processes. Today the primary challenge for adopting something like DAP for a given use case is that there may not be a VDAF that is suitable for that use case. Thus, one way OSTP can support this work is to describe to the research community the government's current set of use cases:

- Which agencies, state or federal, benefit from collecting information about visitors to their websites or other Internet properties? How is this information used? Which statistics or other aggregate artifacts are most useful?
- Are there situations in which information is not collected because it is deemed too sensitive to collect?

Research Opportunities, Risks, and Limitations of PETs (RFI topics #2 and #7)

The RFI's description of privacy-enhancing technologies presenting numerous likely benefits stemming from secure and privacy-preserving data analysis techniques was spot-on. As a technology company with a substantial commitment to furthering research that helps build a better Internet, Cloudflare looks forward to increased collaboration and innovation across industry and academia as a result of the deployment of PETs.

However, Cloudflare also recognizes that deploying PETs is an exercise not without risk. Deploying PETs means less surface area for surveillance by third parties, including governments. Cloudflare has faced resistance in deploying PETs like DOH that limit access to data from industry players that seek to use that data, as well as from governments—including European governments—that are concerned that it will limit the availability of web browsing data. Some governments have shown active opposition to PETs by blocking Internet access or traffic.[21] Such reactions to PETs make the Internet less open, less secure, and less reliable; they

[18] See Richard Barnes, Christopher Patton, & Phillip Schoppmann, Verifiable Distributed Aggregation Functions, IETF, May 26, 2022, at https://datatracker.ietf.org/doc/draft-irtf-cfrg-vdaf/.
[19] See Internet Research Task Force, Crypto Forum Research Group, at https://irtf.org/cfrg.
[20] See Internet Research Task Force, https://irtf.org/.
[21] Examples abound. China, Russa, and South Korea have all been reported to block websites that deploy ESNI, a type of PET that encrypts Internet metadata. See Kevin Bock, David Fifield, Amir Houmansadr, Dave Levin, et al., Exposing and Circumventing China's Censorship of ESNI, Censorship.ai, Aug. 7, 2020, at https://geneva.cs.umd.edu/posts/china-censors-esni/esni/ (China); Nguyen Phong Hoang, Michalis Polychronakis, & Phillipa Gill, Measuring the Accessibility of Domain

could also have the unintended effect of splintering the global Internet.[22] A national strategy that seeks to encourage the deployment of PETs has to integrate international efforts, building consensus around the need to deploy new protocols and approaches. OSTP should also, as part of its national strategy, explicitly enlist the U.S. CTO Team to encourage the U.S. government to leverage those PETs that will help the U.S. government effectively deliver services in a privacy-preserving manner.

Another perceived risk in some quarters of the use of PETs is that they might make systems less reliable. OSTP could help mitigate this perceived risk by supporting research that examines the reliability of PETs, and in working with industry and academic partners to further education on the appropriate use of PETs.

Privacy-enhancing technologies show great promise, but they also have limitations. One limitation of some PETs is that they (intentionally) make it harder for the data collector to "explore" the data. In this case, technologies like OHTTP[23] (the generalization of ODoH) may be worthy of consideration. More generally, it is likely that different use cases will call for the deployment of different PETs. One size will not fit all.

Regulation of PETs (RFI topics #4 and #5)

"Privacy-enhancing technologies" is a new name for a set of tools, techniques, and practices that have been a focus of research and development for many years. At Cloudflare, we take a proactive approach towards privacy, and we believe privacy is not only about responding to different regulations. Privacy is about building technology that helps customers do a better job protecting their users. It is about minimizing the exposure of sensitive data and protecting user privacy and personal information. It is about helping to build a better Internet. In the regulatory arena, Cloudflare has actively supported efforts to develop a framework for U.S. federal privacy standards, urging policymakers to adopt technology-neutral approaches that allow standards to change and improve as technology does. In the context of privacy-enhancing technologies, we encourage the same approach.

In its work across federal agencies and departments, OSTP plays an important role in ensuring that the executive branch understands the effects of regulation, funding, and use of privacy-enhancing technologies. Any regulation of this area should strive to be responsive to the

---

Name Encryption and Its Impact on Internet Filtering, forthcoming Passive and Active Measurement Conference 2022, Feb. 1, 2022, availabl at https://storage.googleapis.com/pub-tools-public-publication-data/pdf/d1cc0cc9b507a8d9ef71c90c35caa9919b95a077.pdf (Russia); and Caleb Chen, South Korea Expands Internet Censorship to HTTPS With First Countrywide Use of SNI, PIABLOG, Feb. 19, 2019, at filteringhttps://www.privateinternetaccess.com/blog/south-korea-expands-internet-censorship-to-https-with-first-countrywide-use-of-sni-filtering/ (South Korea).

[22] See A Declaration for the Future of the Internet, The White House, Apr. 28, 2002, available at https://www.whitehouse.gov/wp-content/uploads/2022/04/Declaration-for-the-Future-for-the-Internet_Launch-Event-Signing-Version_FINAL.pdf.

[23] See Martin Thomson & Christopher A. Wood, Oblivious HTTP, IETF, Feb. 15, 2022, at https://datatracker.ietf.org/doc/draft-ietf-ohai-ohttp/.

dynamic nature of the technology sector, and aim for non-prescriptive and flexible guidance. Any regulation would need to be compatible with other privacy-protecting regulations around the globe, including GDPR, CCPA, and other industry-specific privacy laws in the United States, such as HIPAA. Many of these regulations predate the development of some PETs, and as in any dynamic area of research and development, it is certain that the technology will outpace any regulation.

OSTP, especially its CTO Team, could contribute to the adoption of PETs by incentivizing federal agencies and departments to consider and deploy PETs when reasonable and appropriate to do so. Research into incentive structures or safe harbors for the private sector (including critical infrastructure) around the use of PETs could be informative for regulators worldwide.

Lastly, we believe that strong encryption is key to privacy and security on the Internet, and to the successful development of PETs. OSTP's continued advocacy for strong encryption (and opposition to efforts to weaken encryption), including its use by publicly funded organizations, will go far in improving the privacy of sensitive data.

Recommendations for Advancing and Supporting PETs (RFI topics #1 and #6)

Throughout this response, we have made a number of recommendations for ways that OSTP could advance and support privacy-enhancing technologies. In this section we consolidate them into a single list:

- Across federal agencies and departments, push forward standards-based approaches towards privacy-preserving data sharing and analytics (perhaps, for example, in cooperation with NIST and its Privacy Framework[24]);
- Conduct a baseline assessment of federally funded research on PETs to date and make that baseline assessment publicly available;
- Encourage federal funding of research in PETs (including research that examines the reliability of PETs);
- Work to ensure that any regulation of this area is responsive to the dynamic nature of the technology sector, and aim for non-prescriptive and flexible guidance that is compatible with other privacy-protecting regulations around the globe;
- As part of its national strategy, explicitly enlist the U.S. CTO Team in working to ensure that the U.S. government leverages those PETs that will help the U.S. government effectively deliver services in a privacy-preserving manner;
- Describe to the research community the government's current set of use cases:

[24] See NIST Privacy Framework, at https://www.nist.gov/privacy-framework.

    - Which agencies, state or federal, benefit from collecting information about visitors to their websites? How is this information used? Which statistics or other aggregate artifacts are most useful?
    - Are there situations in which information is not collected because it is deemed too sensitive to collect?

- Integrate international efforts, building consensus around the need to deploy new protocols and approaches to encourage the deployment of PETs;
- Work with industry and academic partners to further education on the appropriate use of PETs; and
- Continue to support the adoption of encryption to protect privacy and oppose efforts to weaken encryption.

Conclusion

Over the next few years, the global tension between the need to keep data as private as possible when using digital services and the perceived need by governments to exert more control over content flowing over the Internet is likely to continue to increase as data becomes increasingly private and encrypted through the development and deployment of privacy-enhancing technologies. Keeping this tension in mind (and supporting research specifically on exploring this tension) will help OSTP formulate a national strategy on privacy-preserving data sharing and analytics that will be more resilient, and resulting policy initiatives that will be productively geared toward responsible use of PETs to benefit individuals and society.

At Cloudflare, we are excited by the promise of privacy-enhancing technologies to protect sensitive information while enabling more innovative and collaborative research. Cloudflare appreciates OSTP's specific and comprehensive questions, and the breadth of expert opinion they will generate for the field. We look forward to continuing to engage with OSTP on privacy-enhancing technologies as this process of developing a national strategy on PETs moves forward.

Sincerely,

/s/ Alissa Starzak

Vice President, Global Head of Public Policy

Attachment: References

References

Selected publications with Cloudflare authors relevant to privacy-enhancing technologies

RFC 9230: Oblivious DNS over HTTPS, Internet Engineering Task Force (IETF). 2022. Eric Kinnear, Patrick McManus, Tommy Pauly, Tanya Verma, Christopher A. Wood.

Might I Get Pwned: A Second Generation Compromised Credential Checking Service, 31th USENIX Security Symposium (USENIX Security 22). Bijeeta Pal, Mazharul Islam, Marina Sanusi, Nick Sullivan, Luke Valenta, Tara Whalen, Christopher A. Wood, Thomas Ristenpart, Rahul Chattejee.

Oblivious DNS over HTTPS (ODoH): A Practical Privacy Enhancement to DNS, Proceedings on Privacy Enhancing Technologies 2021, Volume 4, pp. 575–592. 2021. Sudheesh Singanamalla, Pop Chunhapanya, Jonathan Hoyland, Marek Vavruša, Tanya Verma, Peter Wu, Marwan Fayed, Kurtis Heimerl, Nick Sullivan, Christopher A. Wood.

Privacy Pass: Bypassing Internet Challenges Anonymously, Proceedings on Privacy Enhancing Technologies, no. 3 (2018), pp. 164-180. 2018. Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, Filippo Valsorda.

Selected Cloudflare blog posts relevant to privacy-enhancing technologies

“Empowering Your Privacy,” Emily Hancock, Chief Privacy Officer, 28 January 2020, https://blog.cloudflare.com/empowering-your-privacy/..

“Supporting the latest version of the Privacy Pass Protocol,” Alex Davidson, 28 October 2019, http://blog.cloudflare.com/supporting-the-latest-version-of-the-privacy-pass-protocol/.

“Certifying our Commitment to Your Right to Information Privacy,” Emily Hancock and Rory Malone, 29 July 2021, http://blog.cloudflare.com/certifying-our-commitment-to-your-right-to-information-privacy/.

“Improving DNS Privacy with Oblivious DoH in 1.1.1.1,” Tanya Verma and Sudheesh Singanamalla, 8 December 2020, http://blog.cloudflare.com/oblivious-dns/.

“Cloudflare supports Privacy Pass,” Nick Sullivan, 9 November 2017, http://blog.cloudflare.com/cloudflare-supports-privacy-pass/.

“Cloudflare Zaraz launches new privacy features in response to French CNIL standards,” Yair Dovrat and Yo’av Moshe, 15 June 2022, https://blog.cloudflare.com/zaraz-privacy-features-in-response-to-cnil/.