

Workgroup: Network Working Group

Internet-Draft:

draft-ip-address-privacy-considerations-03

Published: 10 January 2022

Intended Status: Informational

Expires: 14 July 2022

Authors: M. Finkel B. Lassey L. Iannone J.B. Chen
 The Tor Project Google Huawei Google

IP Address Privacy Considerations

Abstract

This document provides an overview of privacy considerations related to user IP addresses. It includes an analysis of some current use cases for tracking of user IP addresses, mainly in the context of anti-abuse. It discusses the privacy issues associated with such tracking and provides input on mechanisms to improve the privacy of this existing model. It then captures requirements for proposed 'replacement signals' for IP addresses from this analysis. In addition, existing and under-development techniques are evaluated for fulfilling these requirements.

Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the mailing list (), which is archived at .

Source for this draft and an issue tracker can be found at <https://github.com/ShivanKaul/draft-ip-address-privacy>.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 14 July 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Terminology](#)
 - [2.1. Categories of Interaction](#)
- [3. IP address tracking](#)
 - [3.1. IP address use cases](#)
 - [3.1.1. Anti-abuse](#)
 - [3.1.2. DDoS and Botnets](#)
 - [3.1.3. Multi-platform threat models](#)
 - [3.1.4. Rough Geolocation](#)
 - [3.2. Implications of IP addresses](#)
 - [3.2.1. Next-User Implications](#)
 - [3.2.2. Privacy Implications](#)
 - [3.3. IP Privacy Protection and Law](#)
 - [3.4. Mitigations for IP address tracking](#)
- [4. Replacement signals for IP addresses](#)
 - [4.1. Signals](#)
 - [4.1.1. Adoption](#)
 - [4.1.2. Privacy Considerations](#)
 - [4.1.3. Provenance](#)
 - [4.1.4. Applying Appropriate Signals](#)
 - [4.2. Evaluation of existing technologies](#)
- [5. Security Considerations](#)
- [6. IANA Considerations](#)
- [7. References](#)
 - [7.1. Normative References](#)
 - [7.2. Informative References](#)
- [Acknowledgments](#)
- [Authors' Addresses](#)

1. Introduction

The initial intention of this draft is to capture an overview of the problem space and research on proposed solutions concerning privacy considerations related to user IP addresses (informally, IP privacy). The draft is likely to evolve significantly over time and may well split into multiple drafts as content is added.

Tracking of IP addresses is common place on the Internet today, and is particularly widely used in the context of anti-abuse, e.g. anti-fraud, DDoS management, and child protection activities. IP addresses are currently used in determining "reputation" [RFC5782] in conjunction with other signals to protect against malicious traffic, since these addresses are usually a relatively stable identifier of a request's origin. Servers use these reputations in determining whether or not a given packet, connection, or flow likely corresponds to malicious traffic. In addition, IP addresses are used in investigating past events and attributing responsibility.

However, identifying the activity of users based on IP addresses has clear privacy implications ([WEBTRACKING1], [WEBTRACKING2]), e.g. user fingerprinting and cross-site identity linking. Many technologies exist today that allow users to obfuscate their external IP address to avoid such tracking, e.g. VPNs ([VPNCMP1], [VPNCMP2]) and Tor ([TOR], [VPNTOR]). Several new technologies are emerging, as well, in the landscape, e.g. Apple iCloud Private Relay [APPLEPRIV], Gnatcatcher [GNATCATCHER], and Oblivious technologies (ODOH [I-D.pauly-dprive-oblivious-doh], OHTTP [I-D.thomson-ohai-ohttp]).

General consideration about privacy for Internet protocols can be found in [RFC6973]. This document builds upon [RFC6973] and more specifically attempts to capture the following aspects of the tension between valid use cases for user identification and the related privacy concerns, including:

- *An analysis of the current use cases, attempting to categorize/group such use cases where commonalities exist.
- *Find ways to enhance the privacy of existing uses of IP addresses.
- *Generating requirements for proposed 'replacement signals' from this analysis (these could be different for each category/group of use cases).
- *Research to evaluate existing technologies or propose new mechanisms for such signals.

With the goal of replacing IP addresses as a fundamental signal, the following sections enumerate existing use cases and describe

applicable substitution signals. This description may not be exhaustive due to the breadth of IP address usage.

2. Terminology

(Work in progress)

This section defines basic terms used in this document, with references to pre-existing definitions as appropriate. As in [RFC4949] and [RFC6973], each entry is preceded by a dollar sign (\$) and a space for automated searching.

*\$ Identity: Extending [RFC6973], an individual's attributes may only identify an individual up to an anonymity set within a given context.

*\$ Reputation: A random variable with some distribution. A reputation can either be "bad" or "good" with some probability according to the distribution.

*\$ Reputation context: The context in which a given reputation applies.

*\$ Reputation proof: A non-interactive zero knowledge proof of a reputation signal.

*\$ Reputation signal: A representative of a reputation.

*\$ Service provider: An entity that provides a service on the Internet; examples services include hosted e-mail, e-commerce sites, and cloud computing platforms.

2.1. Categories of Interaction

Interactions between parties on the Internet may be classified into one (or more) of three categories:

*\$ Private Interaction: An interaction occurring between mutually consenting parties, with a mutual expectation of privacy.

*\$ Public Interaction: An interaction occurring between multiple parties that are not engaged in a Private Interaction.

*\$ Consumption: An interaction where one party primarily receives information from other parties.

3. IP address tracking

3.1. IP address use cases

3.1.1. Anti-abuse

IP addresses are a passive identifier used in defensive operations. They allow correlating requests, attribution, and recognizing numerous attacks, including:

- *account takeover
- *advertising fraud (e.g., click-fraud)
- *disinformation operations (e.g., detecting scaled and/or coordinated attacks)
- *financial fraud (e.g., stolen credit cards, email account compromise)
- *malware/ransomware (e.g., detecting C2 connections)
- *phishing
- *real-world harm (e.g., child abuse)
- *scraping (e.g., e-commerce, search)
- *spam (e.g., email, comments)
- *vulnerability exploitation (e.g., "hacking")

Malicious activity recognized by one service provider may be shared with other services [[RFC5782](#)] as a way of limiting harm.

3.1.2. DDoS and Botnets

Cyber-attackers can leverage the good reputation of an IP address to carry out specific attacks that wouldn't work otherwise. Main examples are Distributed Denial of Service (DDoS) attacks carried out by spoofing a trusted (i.e., having good reputation) IP address (which may or may not be the victim of the attack) so that the servers used to generate the DDoS traffic actually respond to the attackers trigger (i.e., spoofed packets). Similarly botnets may use spoofed addresses in order to gain access and attack services that otherwise would not be reachable.

3.1.3. Multi-platform threat models

As siloed (single-platform) abuse defenses improve, abusers have moved to multi-platform threat models. For example, a public discussion platform with a culture of anonymity may redirect traffic to YouTube as a video library, bypassing YouTube defenses that otherwise reduce exposure of potentially harmful content. Similarly, a minor could be solicited by an adult impersonating a child on a popular social media platform, then redirected to a smaller, less established and less defended platform where illegal activity could occur. Phishing attacks are also common. There are many such cross-platform abuse models and they cause significant public harm. IP addresses are commonly used to investigate, understand and communicate these cross-platform threats. There are very few alternatives for cross-platform signals.

3.1.4. Rough Geolocation

A rough geolocation can be inferred from a client's IP address, which is commonly known as either IP-Geo or Geo-IP. This information can have several useful implications. When abuse extends beyond attacks in the digital space, IP addresses may help identify the physical location of real-world harm, such as child exploitation.

3.1.4.1. Legal compliance

Legal and regulatory compliance often needs to take the jurisdiction of the client into account. This is especially important in cases where regulations are mutually contradictory (i.e. there is no way to be in legal compliance universally). Because Geo-IP is often bound to the IP addresses a given ISP uses, and ISPs tend to operate within national borders, Geo-IP tends to be a good fit for server operators to comply with local laws and regulations

3.1.4.2. Contractual obligations

Similar to legal compliance, some content and media has licensing terms that are valid only for certain locations. The rough geolocation derived from IP addresses allow this content to be hosted on the web.

3.1.4.3. Locally relevant content

Rough geolocation can also be useful to tailor content to the client's location simply to improve their experience. A search for "coffee shop" can include results of coffee shops within reasonable travel distance from a user rather than generic information about coffee shops, a merchant's website could show brick and mortar stores near the user and a news site can surface locally relevant news

stories that wouldn't be as interesting to visitors from other locations.

3.2. Implications of IP addresses

3.2.1. Next-User Implications

When an attacker uses IP addresses with "good" reputations, the collateral damage poses a serious risk to legitimate service providers, developers, and end users. IP addresses may become associated with a "bad" reputation from temporal abuse, and legitimate users may be affected by blocklists as a result. This unintended impact may hurt the reputation of a service or an end user [[RFC6269](#)].

3.2.2. Privacy Implications

IP addresses are sent in the clear throughout the packet journey over the Internet. As such, any observer along the path can pick it up and use it for various tracking purposes. Beside basic information about the network or the device, it is possible to associate an IP address to an end user, hence, the relevance of IP addresses for user privacy. A very short list of information about user, device, and network that can be obtained via the IP address.

- *Determine who owns and operates the network. Searching the WHOIS database using an IP address can provide a range of information about the organization to which the address is assigned, including a name, phone number, and civic address;

- *Through a reverse DNS lookup and/or traceroute the computer name can be obtained, which often contains clues to logical and physical location;

- *Geo-localisation of the device (hence the user) through various techniques [[GEOIP](#)]. Depending on the lookup tool used, this could include country, region/state, city, latitude/longitude, telephone area code and a location-specific map;

- *Search the Internet using the IP address or computer names. The results of these searches might reveal peer-to-peer (P2P) activities (e.g., file sharing), records in web server log files, or glimpses of the individual's web activities (e.g., Wikipedia edits). These bits of individuals' online history may reveal their political inclinations, state of health, sexuality, religious sentiments and a range of other personal characteristics, preoccupations and individual interests;

*Seek information on any e-mail addresses used from a particular IP address which, in turn, could be the subject of further requests for subscriber information.

3.3. IP Privacy Protection and Law

This section aim at providing some basic information about main example of laws adopted worldwide and related to IP address privacy (usually these laws area by product of the broader user privacy protection).

Possible content (to focus only on technical IP address related aspects):

*GDPR (General Data Protection Regulation) - EUROPE: Europe considers IP addresses as personal identification information that should be treated like any other personal information e.g. social security number.

*The United States has opted for a different approach to data protection. Instead of formulating one all-encompassing regulation such as the EU's GDPR, the US chose to implement sector-specific privacy and data protection regulations that work together with state laws to safeguard American citizens' data.

*In 2020, China released the first draft of Personal Information Protection Law (PIPL). The PIPL is the equivalent of European GDPR and will have significant influence.

*Japan Protection of Personal Information (APPI) Act (recent changes put the act close to the GDPR model).

3.4. Mitigations for IP address tracking

The ability to track individual people by IP address has been well understood for decades. Commercial VPNs and Tor are the most common methods of mitigating IP address-based tracking.

*Commerical VPNs offer a layer of indirection between the user and the destination, however if the VPN endpoint's IP address is static then this simply substitutes one address for another. In addition, commercial VPNs replace tracking across sites with a single company that may track their users' activities.

*Tor is another mitigation option due to its dynamic path selection and distributed network of relays, however its current design suffers from degraded performance. In addition, correct application integration is difficult and not common.

*Address anonymization (e.g. [[GNATCATCHER](#)] and similar):

- [\[GNATCATCHER\]](#) is a single-hop proxy system providing more protection against third-party tracking than a traditional commercial VPN. However, its design maintains the industry-standard reliance on IP addresses for anti-abuse purposes and it provides near backwards compatibility for select services that submit to periodic audits.
- [\[APPLEPRIV\]](#) iCloud Private Relay is described as using two proxies between the client and server, and it would provide a level of protection somewhere between a commercial VPN and Tor.

*Recent interest has resulted in new protocols such as Oblivious DNS ([ODoH](#)) and Oblivious HTTP ([OHTTP](#)). While they both prevent tracking by individual parties, they are not intended for the general-purpose web browsing use case.

*Temporary addresses

4. Replacement signals for IP addresses

Fundamentally, the current ecosystem operates by making the immediate peer of a connection accountable for bad traffic, rather than the source of the traffic itself. This is problematic because in some network architectures the peer node of the connection is simply routing traffic for other clients, and any client's use of that node may be only temporary. Ideally, clients could present appropriate identification end-to-end that is separate from the IP address, and uniquely bound to a given connection.

4.1. Signals

There are 7 classes of signals identified in this document that may be used in place of IP addresses. A signal's provenance is a critical property and will be discussed in [Section 4.1.3](#).

*ADDRESS_ESCROW: Provides sufficient information for retroactively obtaining a client's IP address.

*IDENTITY_TRANSPARENCY: Reveals a person's identity within a context.

*IS_HUMAN: Informs the recipient that, most likely, a human recently proved their presence on the opposite end of the connection.

*PEER_INTEGRITY: Provides a secure, remote attestation of hardware and/or software state.

*REIDENTIFICATION: Provides a mechanism for identifying the same user across different connections within a time period.

*REPUTATION: Provides the recipient with a proof of reputation from a reputation provider.

*SOURCE_ASN: Reveals the ASN from which the client is connecting.

In some situations one of the above signals may be a sufficient replacement signal in isolation, or more than one signal may be needed in combination.

Separately, there are three signal categories that are out-of-scope for this document but are important improvements for mitigating abuse on platforms.

*publisher norms: Standard expectations of publishers including identity transparency and conflicts of interest.

*protocol improvements: Increasing security of existing protocols.

*ecosystem improvements: Reducing reliance on less secure systems, for example, migrating user authentication from password-based to WebAuthn [[WEBAUTHN](#)] and relying on multiple factors (MFA).

4.1.1. Adoption

Adoption of replacement signals requires coordination between user agents, service providers, and proxy services. Some user agents and proxy services may support only a subset of these signals, while service providers may require additional signals. A mechanism of negotiation may be needed for communicating these requirements.

In addition, service providers should only require a signal within the scope it will be used. In the same way that service providers only require user authentication when the user requests access to a non-public resource, a signal should not be pre-emptively requested before it is needed. The categories of interaction described above may help define scopes within a service, and they may help communicate to the user the reasoning for requiring a signal.

4.1.2. Privacy Considerations

A signal should not be required without clear justification, service providers should practice data minimization [[RFC6973](#)] wherever possible. Requiring excessive signals may be more harmful to user privacy than requiring IP address transparency. This section provides a more detailed analysis of some signals.

ADDRESS_ESCROW gives service providers a time period within which they may obtain the client's IP address, but the information-in-escrow is not immediately available. Service providers should not gain access to the information in secret. A service provider may misuse the information-in-escrow for tracking and privacy-invasion purposes.

PEER_INTEGRITY partitions users into two groups with valid and invalid hardware/software state, at a minimum. If the signal reveals more information, then it may allow more granular tracking of small sets of devices.

IDENTITY_TRANSPARENCY may expose significant information about a user to a service provider; the resulting privacy invasion may be significantly worse than IP address transparency causes.

IS_HUMAN depends on the mechanism used for proving humanness.

REIDENTIFICATION explicitly allows a service provider to associate requests across unlinkable connections. This signal allows for profiling user behavior and tracking user activity without requesting more identifying information. First-party reidentification is a use case for this signal.

REPUTATION partitions users into a set based on their reputation. The privacy invasion associated with this signal is intentionally small.

SOURCE_ASN allows for identifying request patterns originating from an ASN without providing IP address transparency. However, ASNs are not guaranteed to serve large populations, therefore revealing the source ASN of a request may reveal more information about the user than intended.

4.1.3. Provenance

Replacement signals are only useful if they are trustworthy.

[[OPEN ISSUE: <https://github.com/ShivanKaul/draft-ip-address-privacy/issues/24>]]

4.1.4. Applying Appropriate Signals

As previous discussed, IP addresses are used for various reasons; therefore, describing a one-size-fits-all replacement signal is not appropriate. In addition, the quality and quantity of replacement signals needed by a service depends on the category of interaction of its users and potential attacks on the service.

As an example, the attacks listed above in [Section 3.1.1](#) can be organized into six groups based on the signals that may sufficiently replace IP addresses:

1. IS_HUMAN, REPUTATION, REIDENTIFICATION, PEER_INTEGRITY
 - *advertising fraud (e.g., click-fraud)
 - *phishing
 - *scraping (e.g., e-commerce, search)
 - *spam (e.g., email, comments)
2. IS_HUMAN, REPUTATION, REIDENTIFICATION, ecosystem improvements
 - *account takeover
3. IS_HUMAN, REPUTATION, SOURCE_ASN
 - *influence (e.g., brigading, astroturfing)
4. publisher norms, (publisher) IDENTITY_TRANSPARENCY, PEER_INTEGRITY
 - *disinformation operations (e.g., detecting scaled and/or coordinated attacks)
5. publisher norms, (publisher) IDENTITY_TRANSPARENCY, ADDRESS_ESCROW
 - *real-world harm (e.g., child abuse)
6. IDENTITY_TRANSPARENCY, protocol improvements
 - *financial fraud (e.g., stolen credit cards, email account compromise)

The remaining two attack categories fall outside of the scope of this document.

- *malware/ransomware (e.g., detecting C2 connections)
- *vulnerability exploitation (e.g., "hacking")

Note, IP addresses do not provide a perfect signal in their existing usage, and the above replacement signals do not provide a better signal in all cases.

4.2. Evaluation of existing technologies

Technologies exist that are designed to solve some of the problems described in this document.

Privacy Pass [[I-D.ietf-privacypass-protocol](#)] is a useful building block for solving numerous problems. Its design involves an interaction between a client and server where, at the end, the client is issued a set of anonymous tokens. These tokens may be redeemed at a later time, and this redemption should not be linkable with the initial issuance interaction. One existing use case is substituting a CAPTCHA challenge with a token, where successfully solving a CAPTCHA challenge results in a client being issued a set of anonymous tokens, and these tokens may be used in the future to bypass solving another CAPTCHA challenge. Therefore, Privacy Pass may be acceptable as an IS_HUMAN signal by some service providers. The current token design can't carry additional metadata like a user's reputation or an expiration date, and the tokens are not bound to an identity. The unlinkability property of the tokens is dependent on the implementation of key consistency [[I-D.wood-key-consistency](#)].

Trust Token [[TRUSTTOKEN](#)] is an extension of Privacy Pass where the issuance and redemption functionality are provided in the browser setting. The tokens are allowed to carry public and private metadata as extensions.

Private Access Tokens [[I-D.private-access-tokens](#)] provide a technique for partitioning clients based on a per-origin policy within a time period. Its use cases include rate-limiting access to content and geo-location. PATs could be used as a REIDENTIFICATION signal or a replacement signal for GeoIP, depending on requirements.

5. Security Considerations

This draft discusses IP address use cases, underlying requirements, and possible replacement signals. Adoption challenges and privacy considerations for those signals are also discussed. Further work is needed to build and evaluate these signals as suitable replacements for IP addresses.

6. IANA Considerations

This document has no IANA actions.

7. References

7.1. Normative References

[RFC4949]

Shirey, R., "Internet Security Glossary, Version 2", FYI 36, RFC 4949, DOI 10.17487/RFC4949, August 2007, <<https://www.rfc-editor.org/rfc/rfc4949>>.

[RFC5782] Levine, J., "DNS Blacklists and Whitelists", RFC 5782, DOI 10.17487/RFC5782, February 2010, <<https://www.rfc-editor.org/rfc/rfc5782>>.

[RFC6269] Ford, M., Ed., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, DOI 10.17487/RFC6269, June 2011, <<https://www.rfc-editor.org/rfc/rfc6269>>.

[RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", RFC 6973, DOI 10.17487/RFC6973, July 2013, <<https://www.rfc-editor.org/rfc/rfc6973>>.

7.2. Informative References

[APPLEPRIV] "Apple iCloud Private Relay", n.d., <<https://appleinsider.com/articles/21/06/10/how-apple-icloud-private-relay-works>>.

[GEOIP] Dan, O., Parikh, V., and B. Davison, "IP Geolocation Using Traceroute Location Propagation and IP Range Location Interpolation", Companion Proceedings of the Web Conference 2021, DOI 10.1145/3442442.3451888, April 2021, <<https://doi.org/10.1145/3442442.3451888>>.

[GNATCATCHER] "Global Network Address Translation Combined with Audited and Trusted CDN or HTTP-Proxy Eliminating Reidentification", n.d., <<https://github.com/bslassey/ip-blindness>>.

[I-D.ietf-privacypass-protocol] Celi, S., Davidson, A., and A. Faz-Hernandez, "Privacy Pass Protocol Specification", Work in Progress, Internet-Draft, draft-ietf-privacypass-protocol-01, 22 February 2021, <<https://datatracker.ietf.org/doc/html/draft-ietf-privacypass-protocol-01>>.

[I-D.pauly-dprive-oblivious-doh] Kinnear, E., McManus, P., Pauly, T., Verma, T., and C. A. Wood, "Oblivious DNS Over HTTPS", Work in Progress, Internet-Draft, draft-pauly-dprive-oblivious-doh-09, 5 January 2022, <<https://datatracker.ietf.org/doc/html/draft-pauly-dprive-oblivious-doh-09>>.

[I-D.private-access-tokens]

Hendrickson, S., Iyengar, J., Pauly, T., Valdez, S., and C. A. Wood, "Private Access Tokens", Work in Progress, Internet-Draft, draft-private-access-tokens-01, 25 October 2021, <<https://datatracker.ietf.org/doc/html/draft-private-access-tokens-01>>.

[I-D.thomson-ohai-ohttp] Thomson, M. and C. A. Wood, "Oblivious HTTP", Work in Progress, Internet-Draft, draft-thomson-ohai-ohttp-00, 25 October 2021, <<https://datatracker.ietf.org/doc/html/draft-thomson-ohai-ohttp-00>>.

[I-D.wood-key-consistency] Davidson, A., Finkel, M., Thomson, M., and C. A. Wood, "Key Consistency and Discovery", Work in Progress, Internet-Draft, draft-wood-key-consistency-01, 19 August 2021, <<https://datatracker.ietf.org/doc/html/draft-wood-key-consistency-01>>.

[TOR] "The Tor Project", n.d., <<https://www.torproject.org/>>.

[TRUSTTOKEN] "Trust Token API Explainer", n.d., <<https://github.com/WICG/trust-token-api>>.

[VPNCMP1] Osswald, L., Haeberle, M., and M. Menth, "Performance Comparison of VPN Solutions", Universität Tübingen article, DOI 10.15496/PUBLIKATION-41810, May 2020, <<https://doi.org/10.15496/PUBLIKATION-41810>>.

[VPNCMP2] Khanvilkar, S. and A. Khokhar, "Virtual private networks: an overview with performance evaluation", IEEE Communications Magazine Vol. 42, pp. 146-154, DOI 10.1109/mcom.2004.1341273, October 2004, <<https://doi.org/10.1109/mcom.2004.1341273>>.

[VPNTOR] Ramadhani, E., "Anonymity communication VPN and Tor: A comparative study", n.d., <[Journal of Physics Conference Series](#)>.

[WEBAUTHN] "Web Authentication: An API for accessing Public Key Credentials Level 2", n.d., <<https://www.w3.org/TR/webauthn-2/>>.

[WEBTRACKING1] Bujlow, T., Carela-Espanol, V., Lee, B., and P. Barlet-Ros, "A Survey on Web Tracking: Mechanisms, Implications, and Defenses", Proceedings of the IEEE Vol.

105, pp. 1476-1510, DOI 10.1109/jproc.2016.2637878, August 2017, <<https://doi.org/10.1109/jproc.2016.2637878>>.

[WEBTRACKING2] Mishra, V., Laperdrix, P., Vastel, A., Rudametkin, W., Rouvoy, R., and M. Lopatka, "Don't Count Me Out: On the Relevance of IP Address in the Tracking Ecosystem", Proceedings of The Web Conference 2020, DOI 10.1145/3366423.3380161, April 2020, <<https://doi.org/10.1145/3366423.3380161>>.

Acknowledgments

[[OPEN ISSUE: TODO]]

Authors' Addresses

Matthew Finkel
The Tor Project

Email: sysrqb@torproject.org

Bradford Lassey
Google

Email: lassey@chromium.org

Luigi Iannone
Huawei Technologies France S.A.S.U

Email: luigi.iannone@huawei.com

J. Bradley Chen
Google

Email: bradchen@google.com