Rethinking Geolocalization on the Internet

Augustin Laouar[1], Loïc Desgeorges[2], Paul Schmitt[3], and Francesco Bronzino[1,4]
[1]ENS Lyon, [2]UCBL, [3]Cal Poly and ICSI, [4]Institut Universitaire de France

Abstract

Location underpins critical Internet services, yet our primary mechanism for Internet localization, IP-based geolocation, fails to meet the needs of all stakeholders. User location is conflated with network location, leading to a fundamental mismatch between the goals of content providers, infrastructure operators, and regulators. As users increasingly adopt privacy-preserving technologies that obscure their network identity, this mismatch becomes more pronounced, making localization even more challenging. This paper argues that the problem cannot be solved by simply improving the accuracy of incumbent mechanisms that are inappropriately applied today to solve multiple, unrelated problems. Instead, we require a new approach for localization on the Internet.

CCS Concepts

• **Networks** → **Network architectures**; • **Security and privacy** → **Privacy protections**.

Keywords

Geolocalization, Internet Architecture, Privacy, Network Measurement

ACM Reference Format:
Augustin Laouar[1], Loïc Desgeorges[2], Paul Schmitt[3], and Francesco Bronzino[1,4] . 2025. Rethinking Geolocalization on the Internet. In *The 24th ACM Workshop on Hot Topics in Networks (HotNets '25), November 17–18, 2025, College Park, MD, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3772356.3772421

1 Introduction

> I remember seeing an elaborate and complicated automatic washing machine for automobiles that did a beautiful job of washing them. But it could do only that, and everything else that got into its clutches was treated as if it were an automobile to be washed. I suppose it is tempting, if the only tool you have is a hammer, to treat everything as if it were a nail.
> *The Psychology of Science*, Maslow

Location plays a critical role on the Internet, shaping how we interact with digital services and infrastructure. From content delivery networks that optimize routing based on proximity to infrastructure, to content restrictions that vary based on region, or security systems that identify suspicious behavior by detecting unusual locations, spatial awareness is woven into the fabric of online systems.

Despite this fundamental importance, our current approach to Internet localization is deeply flawed. The dominant method (*i.e.*, IP geolocation) assumes that network identifiers (IP addresses designed for routing packets) can reliably map to physical geography. This assumption has always been problematic [20, 29, 33], but it persists largely because it requires no user interaction and functions as a relatively lightweight, if imperfect, solution [12, 38]. Further, the current approach conflates two distinct concepts: the location within a network (*i.e.*, to inform routing decisions focused on infrastructure efficiency) and the geographic position of the user for policy and content decisions. This conflation is not just a technical oversight; it inherently misaligns with the problems that the different stakeholders must solve. Content restriction, infrastructure optimization, and regulatory compliance have fundamentally different requirements that cannot be adequately addressed through a single location primitive.

The research community, as well as industry, have tried for years to solve the issue enhancing IP geolocation's accuracy. Unfortunately, the inadequacy of the incumbent approach is exacerbated as users increasingly look to enhance their online privacy by masking their IP addresses. Once confined to tools such as VPNs and Tor, this has now become a standard feature integrated by major platforms: Apple, Google, and Microsoft [2, 13, 24] now offer IP privacy to hundreds of millions of users. These services deliberately break the

HotNets '25, College Park, MD, USA

ACM ISBN 979-8-4007-2280-6/25/11
https://doi.org/10.1145/3772356.3772421

connection between a user and their network identity, rendering IP-based localization even more challenging.

In this work, we argue that the community must rethink how to achieve Internet localization, accounting for the different needs of stakeholders and the growing importance of user privacy. We highlight the limitations of the current approach with a short study analyzing the location accuracy of a commercial IP geolocation service (IPinfo) for IP addresses used by Apple's iCloud Private Relay, a widely adopted privacy-preserving technology. Our measurements reveal that despite attempts to maintain geographic coherence between relay exit nodes and user locations, fundamental tensions remain unresolved.

Given our findings, we sketch a new design that explicitly decouples different location requirements and creates purpose-specific mechanisms for each and we map the research challenges that must be solved before such a system can become practical. Our goal with this work is to foster a broader discussion around Internet localization within the networking community. By reimagining how location is determined and verified online, we look to enable systems that provide increased utility for all stakeholders.

2 Context

We briefly review the current state and challenges of Internet localization.

2.1 Geolocalization Today

The modern Internet relies on IP geolocation as the predominant method for determining the geographic location for both devices and users. At its core, IP geolocation aims to infer a client's position from network-level evidence tied to its public IP address. To achieve this goal, commercial providers (*e.g.*, IPinfo, MaxMind, *etc.*) combine static evidence (RIR allocations, WHOIS, routing tables) with dynamic signals (reverse-DNS lexica, end-host telemetry, and latency triangulation), to place each IP probabilistically on the map [15, 23]. Given its simplicity, ubiquity, and low cost, IP-based geolocation continues to underpin operational workflows that require spatial cues, from data-residency compliance and location-tailored search results to security filtering [6, 38].

Unfortunately, IP geolocation has been long known to be an unreliable, often greatly inaccurate, technique [6, 20, 22, 29, 33]. IP geolocation services assume each public address maps to a single stable place. Yet, large-scale address reuse, anycast content delivery, and policy-driven BGP routing systematically break that premise, pushing the same address to users or replicas that can be hundreds of kilometers apart.

To address this issue, most prior work has focused on enhancing IP geolocation techniques through various approaches: exploiting known network landmarks [36], leveraging systems behavior [30], applying data-driven methods [37], and other refinements [38]. While these efforts improve accuracy in traditional settings, they all share a fundamental limitation: their continued reliance on the public IP address as the primary identifier for geolocation. Unfortunately, this dependency becomes increasingly problematic as the core assumption, that network identifiers can reliably map to physical geography, further breaks down.

2.2 Geolocalization Tomorrow

In recent years, privacy-preserving tools like VPNs and Tor have seen increased adoption as users look for ways to protect their online privacy and circumvent censorship or geographic restrictions imposed by content providers and ISPs. From a geolocation perspective, these tools introduce a significant challenge: by obscuring users' actual IP addresses, they make accurate IP geolocation more difficult. While one could argue that such tools will never achieve mainstream adoption due to their complexity and performance overhead, new and easier-to-use solutions have recently emerged. Privacy-preserving overlay systems now embed anonymization directly within operating systems or browsers. Apple's iCloud Private Relay (PR) [2], Google's IP Protection [13], and Microsoft's Edge Secure Network [24] are examples of architectures that route traffic through multi-hop tunnels built on the MASQUE protocol, forming a performant privacy-preserving overlay network [28].

If and when mainstream systems and browsers enable this feature by default, a substantial portion of Web traffic will traverse these relays, further complicating geolocation tasks. To mitigate this challenge, relay providers attempt to preserve geolocation fidelity by publicly disclosing information about egress IP addresses and their logical locations. For instance, Apple and its infrastructure partners regularly publish egress IP prefixes and their associated locations [3] to inform external services (*i.e.*, IP geolocation providers) of "accurate" Private Relay user locations. However, as we discuss in Section 3, even with such ground-truth (which represents the best case for overlay networks, since most do not expose similar data), user geolocation remains difficult and prone to significant errors.

2.3 What Now?

This evolving landscape leads to two key effects. First, the concept of geolocation is becoming increasingly confusing: services now face the dual task of mapping IP addresses to both infrastructure and user locations, conflating fundamentally different objectives into a single API. For example, with relay services, commercial IP-geolocation databases must decide whether to map the observable infrastructure (*i.e.*, the relay egress node) or the users behind them, who may be located hundreds of kilometers away. This semantic mismatch induces systematic discrepancies that will scale with adoption and, if left unaddressed, threaten accuracy-critical services. Second, most current patches depend on commercial entities whose primary motivations are self-interest rather than the integrity of Internet infrastructure[1]. This results in a fragmented and unreliable ecosystem that is subject to the whims of private companies.

These factors raise a fundamental question: should we continue patching the existing IP geolocation ecosystem, or does the shift toward privacy-preserving browsing represent an opportunity for a clean slate approach? Rather than relying on ad-hoc solutions around an increasingly obsolete paradigm, we argue that the time has come for the networking community to *reimagine geolocation* for an era where IP addresses are no longer de facto identifiers of Internet users.

3 Case Study: Private Relay

In this section, we examine the accuracy of commercial IP geolocation services when used with privacy-preserving overlay systems. Our measurements reveal that despite efforts to maintain geographic coherence between relay exit nodes and user locations, significant challenges persist.

3.1 Geolocation Discrepancy

For our study, we choose Apple's iCloud Private Relay (PR), a browser-integrated relay deployed at large scale (approximately 280,000 egress IPs). PR is a privacy-preserving overlay that provides network-layer privacy by routing user traffic through two relays: the first operated by Apple and the second by a third-party CDN (*e.g.*, Akamai, Cloudflare, or Fastly). This segmented approach ensures that no single entity can correlate both the user's identity and their complete browsing activity, effectively decoupling the direct relationship between client and server while maintaining functional internet connectivity. However, this architecture introduces a new challenge for IP geolocation services: the user's location is no longer directly tied to the egress IP address, as the relay decouples the user's physical position from the egress infrastructure. To compensate for this, Apple publishes a list of mappings between egress IP address ranges and the city where the PR users are located [3].

This list is intended to be consumed by IP geolocation services, allowing them to update their databases with "correct" IP-to-location mappings. In practice, however, discrepancies frequently arise between Apple's published geofeed and the locations reported by commercial geoIP providers—the same databases consulted by location-based services (LBS). We aim to quantify the extent of this discrepancy, as well as its potential impact on users that access location-based services that rely on IP geolocation.

3.2 Global Analysis

We begin our analysis by studying discrepancies between geolocations reported by Apple and those provided by IPinfo [19], a widely used commercial IP geolocation service[2]. We collect Apple's mappings by retrieving the published CSV file containing PR IP prefixes mapped to city, state (or region), and country. To assess the evolution of PR egress IPs over time, we download the file daily from March 22, 2025, to June 22, 2025. We convert Apple's geographic labels into precise coordinates by geocoding each location using both Nominatim [26] (OpenStreetMap) and the Google Geocoding API [14].[3] Similarly, we download the IPinfo database daily and resolve every PR egress IP against the database. For each egress IP, we compute the distance between Apple's coordinates and those returned by the IP geolocation service.

Figure 1 presents the cumulative distribution of geolocation discrepancies, grouped by continent, for both IPv4 and IPv6 addresses (we observe similar results for both versions, thus we aggregate them in a single plot). PR's impact spans all continents: many egress IPs are located tens to hundreds of kilometers away from Apple's declared position, with 5 % exhibiting differences exceeding 530 km. Because Apple operates relays in nearly every country, cross-border misplacements are

[1]While we assume the companies involved do not explicitly wish to degrade Internet localization, their design choices may not generalize.

[2]Certainly, other geolocation services may perform better or worse compared with IPinfo. This study is not meant to be an exhaustive IP geolocation measurement campaign with respect to Private Relay. Our goal is to explore discrepancies such overlays introduce.

[3]When the resulting coordinates differed by less than 50 km, we selected Google's result. For discrepancies exceeding 50 km, we manually verified and selected the more accurate coordinate pair.

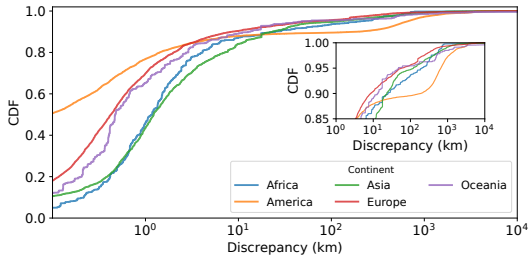

Figure 1: Geolocation discrepancy by continent.

rare—only 0.5 % of egresses are mapped by the IP geolocation service to the wrong country. However, many location-based services require finer-grained accuracy, and differences within a country can have significant consequences–especially in nations where legislation varies by state or province. For instance, state-level mismatches (*i.e.*, when the two services identify different states within a country) affect 11.3 % of egresses in the United States, 9.8 % in Germany, and 22.3 % in Russia. These findings confirm that the distortions introduced by PR are global and structural rather than incidental. We initially hypothesized that the root causes for these differences lay in a temporal delay between updates across databases. However, throughout our measurement campaign, we tracked every egress addition or relocation announced by Apple—fewer than 2,000 events in total. The IP geolocation service consistently reflected these changes with 100 % accuracy, ruling out data staleness as the cause of the mismatches.

3.3 Cause of Discrepancies

Given our findings, we aim to separate *classical* IP geolocation errors (*i.e.*, typical of commercial databases) from mismatches introduced by the use of PR. To identify their origin, we performed a latency-based validation using RIPE Atlas on May 28, 2025 [31]. For discrepancies exceeding 500 km, we selected up to 10 nearby probes for each candidate location and measured RTTs to the IP prefix. These RTTs were used in a temperature-controlled softmax to estimate the most likely location. We focus on the United States, which concentrates most PR egress prefixes (63.7 % on May 28, 2025) and offers dense RIPE Atlas coverage (1,663 active probes). To ensure temporal consistency and limit measurement overhead, we use a single snapshot. Finally, as Apple publishes very large IPv6 prefixes (*i.e.*, /45, /64) that are far too vast for exhaustive probing, a preliminary random sampling inside each prefix showed that geolocation outputs are invariant across addresses. We therefore test only the first two IP addresses of every advertised IPv6 range, whereas for IPv4, we probe all listed addresses.

Table 1: RIPE Atlas validation of > 500 km differences (USA, 28 May 2025).

Outcome	Count	Share (%)
IP geolocation discrepancies	5982	60.12
PR-induced discrepancies	3264	32.80
Inconclusive	704	7.08

Our validation (Table 1) shows that 60.12 % of the > 500 km discrepancies arise because IPinfo mislocates the egress, a pattern attributable to the inherent limitations of classic IP-based geolocation. In contrast, for 32.80 % of cases, the probes agree with IPinfo: the database seems to correctly point to the relay's egress POP, yet Apple reports the user's chosen city.

3.4 IPinfo's Comments

To better understand the discrepancies observed between Apple's iCloud Private Relay (PR) geofeed and IPinfo's reported locations, we contacted IPinfo and shared our findings. Their detailed feedback provided valuable insight into several sources of error in their current geolocation process and its interaction with geofeeds, which they have since corrected.

IPinfo acknowledged that part of the incorrect entries originated from *user-submitted corrections* that had inadvertently overridden accurate data coming from trusted geofeeds. These erroneous updates have since been deleted, and IPinfo indicated they will modify their ingestion pipeline to prevent such corrections from superseding verified sources in the future. They also identified additional mismatches caused by *geocoding errors* within IPinfo's internal pipeline. These issues primarily affected sparsely populated areas and locations referenced by administrative regions (*e.g.*, county or area names) rather than precise settlements. IPinfo reported implementing multiple fixes to improve the handling of such ambiguous cases. IPinfo further noted that part of the discrepancy stemmed from inaccuracies in our own geocoding of Apple's iCloud geofeed entries using Google Maps and Nominatim. According to their assessment, approximately 0.8% of the entries were incorrectly resolved in our dataset, with around 32% of these misplacements exceeding 1,000 km.

Overall, IPinfo emphasized that this analysis highlights the inherent challenges of working with *geofeeds*, due to the absence of standardized and unambiguous geographical identifiers. Even well-formed geofeeds—such

as Apple's iCloud feed, constructed with reliable geocoding services—can produce significant misplacements. IPinfo added that while they identify IPs that are not included in trusted feeds through active measurements (*e.g.*, ping latency), for trusted feeds like Private Relay they intentionally rely on the geofeed's declared coordinates, which remain subject to geocoding uncertainty.

In conclusion, IPinfo's feedback underscores the difficulty of accurately reflecting a trusted geofeed, even when treated as a ground truth. These challenges are expected to be even greater for overlay networks that do not expose verified geographic anchors, such as VPNs or mobile networks.

4 Discussion

As emerging systems increasingly decouple users from their network egress points, traditional IP geolocation approaches face fundamental limitations that cannot be addressed through incremental database improvements alone. In this section, we discuss the path forward, arguing that the research community should take a leading role in addressing this problem by rethinking how geolocation operates on the Internet.

4.1 Whither IP Geolocation?

While case-specific techniques such as those used for Private Relay may help temporarily maintain user-centric geolocation accuracy, they remain ad-hoc and unsustainable. Each new overlay would require its own external mapping and coordination, multiplying dependencies and potential failure points. More fundamentally, other overlay networks—such as commercial VPNs or mobile carrier networks—lack any authoritative source capable of disclosing user locations without undermining privacy. As such, Private Relay represents a convenient but exceptional case where a ground truth exists; the growing diversity of overlay systems makes incremental patching both fragile and unsustainable.

Our measurements demonstrate how discrepancies arise from using a network technology—*i.e.*, IP geolocation—to infer user positions on today's Internet. This approach conflates two fundamentally different objectives: localizing network infrastructure versus localizing users. This confusion becomes increasingly problematic as overlay networks decouple users from their egress points. Nevertheless, IP geolocation remains highly valuable when used for its intended purpose: mapping network infrastructure. CDNs effectively leverage IP geolocation, combined with active measurements such as traceroute and latency probes, as well as passive techniques like BGP route inspection and real-user monitoring, to identify optimal points of presence for content delivery [17, 21]. Network operators use IP geolocation to detect routing anomalies, identify potential security threats based on traffic origin, and optimize peering relationships. Research communities rely on IP geolocation to study Internet topology [34], distributed denial-of-service attacks [35], and global traffic patterns. Overall, IP geolocation excels at its intended purpose and should be used accordingly. Problems arise only when it is repurposed to locate users, a distinction that becomes increasingly critical as overlay networks reshape the Internet landscape. Instead, a new dedicated service should be created for the purpose of localizing *users*.

4.2 User Localization: a Wishlist

Rethinking user localization for the Internet will require meeting several fundamental properties, while also addressing their associated research challenges. We list these properties below, highlighting the inherent trade-offs that must be resolved to create a practical and effective user localization system.

Accuracy. A user localization system should provide accurate and reliable location information to satisfy services that demand precise geographic data. In contrast to current approaches, the concept of accuracy should be clearly defined and quantifiable as distance error relative to an actual user's location (*e.g.*, within 10 km for city-level granularity), rather than the uncertainty that arises from conflating users and infrastructure. However, this property inherently creates trade-offs with the privacy and verification requirements that follow.

Verifiability. The location information obtained should be trusted by the receiving service, which implies that the user's position has been previously verified through reliable mechanisms. This verification process must balance trust with user privacy, potentially requiring lightweight cross-checks such as latency triangulation, BGP consistency, or hardware attestation, while avoiding overly intrusive validation methods.

Privacy-conscious. Following Moore's definition [25], privacy is best understood as a right of individuals to control access to and use of their personal information, by deciding when, how, and to what extent that information is shared with others. Users should then retain control over the granularity of the location information they provide, from country-level to city-level or more precise coordinates depending on the service requirements and their privacy preferences. This granularity

control creates inherent trade-offs with accuracy and verifiability, as coarser location data may be less useful for services while being harder to verify independently.

Scalable. For practical deployment, a localization system should be lightweight enough to handle Internet-scale usage without imposing significant computational or network overhead on users, services, or the network infrastructure in general. This might be challenging if verification mechanisms are implemented for all localization requests, as comprehensive verification could introduce substantial computational and network costs.

Frictionless. As user experience is a key factor on today's Internet, users should not be oversolicited by location requests or verification procedures. The system should operate transparently in the background, minimizing user intervention while respecting their privacy choices. This creates tension with verification requirements, as more robust verification mechanisms may introduce additional friction.

Open. Prior works provide partial building blocks toward a privacy-preserving geolocation system [16, 27], but do not offer a complete solution. To ensure transparency, trust, and broad interoperability, such a system should be open, publicly specified through standardization bodies, and built from the ground up for independent implementation and verification.

4.3 Rethinking User Geolocalization

The properties listed in the previous Section are not always compatible and involve inherent trade-offs. Designing a system that satisfies all these properties while resolving their trade-offs is highly challenging and, to date, largely unresolved. However, we believe that a practical user geolocalization system that meets these properties is possible. As a first step toward this direction, we sketch a high-level design for a localization system. Our proposal builds on the expertise of our community in building trustworthy, scalable distributed systems for various goals. In this design, a trusted third party attests both the user's position (furnished via reliable signals) and the minimum spatial granularity required by the service, as dictated by its functional requirements. Trust among the third party, the user, and a location-based service (LBS) should be anchored in a certificate chain, analogous to the X.509 trust chain that secures Internet connections. To ensure scalability and minimize latency, the third party should operate offline, issuing long-lived certificates that define each LBS's authorization scope and short-lived tokens attesting user positions, without being involved in subsequent connections. One possible

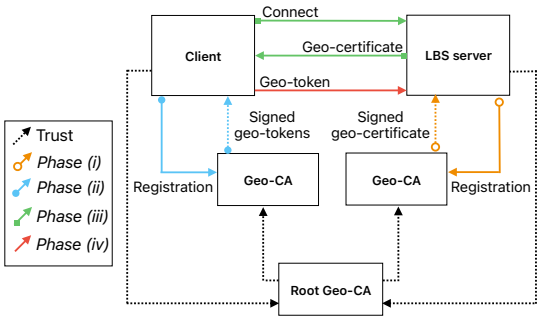

Figure 2: Geo-CA workflow.

design could exchange and verify these certificates and tokens during the TLS handshake between the client (*i.e.*, the software agent representing the user) and the server, thereby integrating localization proofs directly into the secure channel establishment process.

We refer to this system as Geo-Certification Authorities (Geo-CAs), and illustrate its workflow in Figure 2, which unfolds in four high-level phases: **(i) LBS registration.** Each LBS registers with one or more Geo-CAs and obtains a long-lived signed certificate (*e.g.*, one-year validity) that attests the finest spatial granularity it is authorized to request. **(ii) User registration.** The client periodically uploads its position to the selected Geo-CAs and receives a bundle of signed geo-tokens—one per admissible granularity level (*e.g.*, exact point, neighborhood, city, region, country), each embedding the issuer's identity, the user's position, an expiry time, and any extra metadata a service might later require. **(iii) Server authentication.** The server presents its Geo-CA certificate, which attests to the spatial granularity it is authorized to request. The client verifies the certificate chain against its trusted root Geo-CAs. **(iv) Client attestation.** The client sends a geo-token of the requested granularity, which the server verifies to attest the authenticity of the submitted location information. At this point, the client's location is confirmed.

4.4 Open Challenges

Although preliminary, this design sketches a feasible path toward privacy-preserving location attestation and surfaces several open challenges.

Token Replay. Preventing token replay is essential to avoid location fraud, where an attacker reuses a valid geo-token for unauthorized access. Mechanisms such as DPoP [5], which bind tokens to ephemeral public keys and per-session challenges, provide a promising starting point. However, they must be carefully adapted

to prevent linkability across sessions, illustrating the fundamental trade-off between *privacy* and *verifiability*.

Privacy-Preserving Issuance. Similar privacy challenges arise in DNS, which has inspired solutions such as oblivious resolution [18] that separates user identity from query content through split trust between independent entities. Following this principle, Geo-CA architectures could use intermediaries to decouple user identity from attested location, or explore complementary mechanisms such as rotating authorities to further limit information linkage. Moreover, Geo-CAs must not learn users' raw coordinates. Privacy-preserving issuance techniques such as zero-knowledge region proofs [11] and blind signatures [4, 7, 10] can help protect user privacy. Notably, prior work [32] showed that millions of blind signatures can be processed per second with negligible overhead, indicating these methods scale efficiently. However, they complicate lightweight verification of the user's position, such as latency-based checks. This reflects a core tension between *privacy* and *scalability*.

Position Updates. Determining optimal update frequency presents a fundamental challenge. Frequent updates degrade *privacy* (by revealing mobility patterns) and reduce *frictionless operation* (*e.g.*, battery drain on mobile devices, increased network traffic). Conversely, infrequent updates compromise *accuracy*, as tokens become stale for mobile users. A practical system must balance token freshness against overhead, potentially through adaptive strategies that adjust update frequency based on movement or context.

Resilience. Geo-CAs introduce points of failure, where outages could prevent token and certificate issuance. To mitigate this, the system could draw inspiration from DNS, leveraging redundancy, distribution, and failover to ensure availability.

Governance and Regulation. Beyond technical safeguards, governance remains a key concern: we must avoid replicating the centralization risks seen in Web PKI [1, 8]. A more resilient model could rely on *federated trust*. This approach could draw inspiration from existing systems such as Google's Certificate Transparency (CT) [9], which distributes accountability across multiple independent log operators. Combining federated trust with public transparency would reduce single points of control while ensuring verifiable and accountable operation. Furthermore, establishing *open regulatory standards* could define how Geo-CAs determine and enforce the level of spatial granularity each service is authorized to request, based on its legitimate operational needs. Such standards would formalize least-privilege principles for location access, ensuring that only the minimum necessary detail is disclosed while preserving user privacy.

Adoption. While the technical architecture addresses core limitations of current geolocation systems, widespread adoption will remain a central challenge. Success depends on providing compelling incentives for all stakeholders. For *users*, the system must deliver tangible benefits: reliable access to location-based services (avoiding content blocking due to VPN use or inaccurate IP geolocation), user-controlled privacy (granular disclosure rather than all-or-nothing), and frictionless operation (no additional authentication steps beyond normal browsing). For *services*, key incentives include verifiable localization resistant to manipulation through VPNs, reduced liability from regulatory compliance, and potentially lower costs compared to maintaining proprietary geolocation solutions. However, the system introduces implementation overhead and regulatory constraints. Adoption may follow a gradual path: initial deployment for high-stakes use cases (*e.g.*, content licensing, regulated services) where verification benefits outweigh costs, followed by broader adoption as infrastructure matures and browsers integrate native support. Critically, addressing the aforementioned technical and governance challenges will determine whether Geo-CAs can achieve sufficient momentum to become a viable solution for user localization.

5 Conclusion

IP-based geolocation remains the default method for Internet localization. While its methods are well-suited for locating from a network-centric perspective, locating users demands different approach as clients increasingly disconnect from their network egress points.

We propose exploring a distributed Geo-Certification system that acts as a trust intermediary between users and location-based services, attesting both the user's reported location and the level of granularity required by the service. Such a system could improve location accuracy while preserving user privacy, with minimal friction. While significant research challenges remain—both in design and deployment—we believe the networking community is well-positioned to advance this direction. We hope our work encourages further exploration of robust, privacy-preserving geolocation architectures.

Acknowledgement. We thank the reviewers and our shepherd Shaddi Hasan for their constructive feedback. We deeply thank IPinfo for their availability to provide insights related to our results. This work was supported by the ANR-24-CE25-1133 (GTTP) Project at ENS Lyon.

References

[1] G. Akiwate et al. 2024. On the Centralization and Regionalization of the Web. *arXiv preprint arXiv:2405.12345* (2024).

[2] Apple Inc. 2021. iCloud Private Relay — Overview. https://www.apple.com/icloud/docs/iCloud_Private_Relay_Overview_Dec2021.pdf. White paper, accessed 16 Jun 2025.

[3] Apple Inc. 2025. Private Relay Egress IPs. https://mask-api.icloud.com/egress-ip-ranges.csv. CSV file, accessed 20 Jun 2025.

[4] Bellare, Namprempre, Pointcheval, and Semanko. 2003. The one-more-RSA-inversion problems and the security of Chaum's blind signature scheme. *Journal of Cryptology* 16 (2003), 185–215.

[5] Daniel Benjamin, Torsten Lodderstedt, and Daniel Fett. 2023. *OAuth 2.0 Demonstrating Proof-of-Possession at the Application Layer (DPoP).* Technical Report RFC 9449. Internet Engineering Task Force. doi:10.17487/RFC9449

[6] Patricia Callejo, Marco Gramaglia, Ruben Cuevas, and Angel Cuevas. 2022. A deep dive into the accuracy of IP Geolocation Databases and its impact on online advertising. *IEEE Transactions on Mobile Computing* 22, 8 (2022), 4359–4373.

[7] David Chaum. 1983. Blind Signature System.. In *Crypto*, Vol. 83. Springer, 153.

[8] L. Cordova Azevedo et al. 2025. Assessing SSL/TLS Certificate Centralization: Implications for Digital Sovereignty. *arXiv preprint arXiv:2501.01234* (2025).

[9] Al Cutter, Ben Laurie, Eran Messeri, and Adam Langley. 2021. Certificate Transparency Version 2.0. RFC 9162. doi:10.17487/RFC9162 Status: Standards Track.

[10] Alex Davidson, Ian Goldberg, Nick Sullivan, George Tankersley, and Filippo Valsorda. 2018. Privacy pass: Bypassing internet challenges anonymously. *Proceedings on Privacy Enhancing Technologies* (2018).

[11] Jens Ernstberger, Chengru Zhang, Luca Ciprian, Philipp Jovanovic, and Sebastian Steinhorst. 2024. Zero-Knowledge Location Privacy via Accurate Floating-Point SNARKs. *arXiv preprint arXiv:2404.14983* (2024).

[12] Fact.MR. 2024. *IP Geo-Location Market Study by Content Personalization, Fraud Detection, Ad Targeting, Traffic Analysis, Compliance, Geo-Targeting, Geo-Fencing, and Digital Rights Management from 2024 to 2034.* Technical Report. Fact.MR. https://www.factmr.com/report/4703/ip-geo-location-market Market size US$2.74 billion in 2024; accessed 12 Jun 2025.

[13] Google Chrome Team. 2025. IP Protection: Privacy Sandbox Explainer. https://developer.chrome.com/docs/privacy-sandbox/ip-protection/. Accessed 16 Jun. 2025.

[14] Google Maps Platform. 2025. Geocoding API Documentation. https://developers.google.com/maps/documentation/geocoding Accessed: 2025-07-01.

[15] Matthieu Gouel, Kevin Vermeulen, Olivier Fourmaux, Timur Friedman, and Robert Beverly. 2021. A Longitudinal Study of an IP Geolocation Database. *Proceedings of the Traffic Measurement and Analysis Conference (TMA)* (2021). arXiv:2107.03988.

[16] Scott Hendrickson and Christopher A. Wood. 2023. Privacy Pass Geolocation Hint Extension. Internet-Draft draft-hendrickson-privacypass-geo-extension-00. https://www.ietf.org/archive/id/draft-hendrickson-privacypass-geo-extension-00.html Work in Progress.

[17] Peter Hillmann, Lars Stiemert, Gabi Dreo Rodosek, and Oliver Rose. 2015. Modelling of IP Geolocation by use of Latency Measurements. In *2015 11th International Conference on Network and Service Management (CNSM)*. IEEE, 173–177.

[18] Paul E. Hoffman and Eric Nygren. 2022. Oblivious DNS over HTTPS. RFC 9230. doi:10.17487/RFC9230

[19] IPinfo. 2025. IPinfo - IP address data API and geolocation service. https://ipinfo.io. Accessed: 2025-07-09.

[20] Dan Komosny, Miroslav Voznak, and Saeed Ur Rehman. 2017. Location accuracy of commercial IP address geolocation databases. *Information technology and control* 46, 3 (2017), 333–344.

[21] Rupa Krishnan, Harsha V Madhyastha, Sridhar Srinivasan, Sushant Jain, Arvind Krishnamurthy, Thomas Anderson, and Jie Gao. 2009. Moving beyond end-to-end path information to optimize CDN performance. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement*. 190–201.

[22] Ioana Livadariu, Thomas Dreibholz, Anas Saeed Al-Selwi, Haakon Bryhni, Olav Lysne, Steinar Bjørnstad, and Ahmed Elmokashfi. 2020. On the accuracy of country-level IP geolocation. In *Proceedings of the 2020 Applied Networking Research Workshop*. 67–73.

[23] MaxMind Inc. 2025. Choose the IP Intelligence Data You Need. https://support.maxmind.com/hc/en-us/articles/4408200217371-Choose-the-IP-Intelligence-Data-you-Need. Accessed 13 Jun. 2025.

[24] Microsoft Corporation. 2024. Use the Microsoft Edge Secure Network to Protect Your Browsing. https://support.microsoft.com/en-us/topic/use-the-microsoft-edge-secure-network-to-protect-your-browsing-885472e2-7847-4d89-befb-c80d3dda6318. Accessed 16 Jun. 2025.

[25] Adam D Moore. 2008. Defining privacy. *Journal of Social Philosophy* 39, 3 (2008), 411–428.

[26] Nominatim. 2025. Open-source Geocoding with OpenStreetMap Data. https://nominatim.org Accessed: 2025-07-01.

[27] Markus Pauly. 2023. An HTTP Geolocation Hint for IP Addresses. Internet-Draft draft-pauly-httpbis-geoip-hint-01. https://datatracker.ietf.org/doc/html/draft-pauly-httpbis-geoip-hint Work in Progress.

[28] Tommy Pauly, Eric Kinnear, David Schinazi, and Mike Bishop. 2023. Proxying IP in HTTP. RFC 9484, Internet Engineering Task Force. https://www.rfc-editor.org/rfc/rfc9484

[29] Ingmar Poese, Steve Uhlig, Dali Kaafar, Benoit Donnet, and Bamba Gueye. 2011. IP Geolocation Databases: Unreliable? *ACM SIGCOMM Computer Communication Review* 41, 2 (2011), 53–56. doi:10.1145/1971162.1971172

[30] Hugo Rimlinger, Kevin Vermeulen, Olivier Fourmaux, and Timur Friedman. 2024. Poster: GeoResolver, An Accurate, Scalable, and Explainable Geolocation Technique Using DNS Redirection. In *Proceedings of the 20th International Conference on emerging Networking EXperiments and Technologies*. 21–22.

[31] RIPE NCC. 2025. RIPE Atlas. https://atlas.ripe.net/. Accessed June 2025.

[32] Paul Schmitt and Barath Raghavan. 2021. Pretty good phone privacy. In *30th USENIX Security Symposium (USENIX Security 21)*. 1737–1754.

[33] Yuval Shavitt and Noa Zilberman. 2011. A Geolocation Databases Study. *IEEE Journal on Selected Areas in Communications* 29, 10 (2011), 2044–2056. doi:10.1109/JSAC.2011.111209

[34] Ye Tian, Ratan Dey, Yong Liu, and Keith W Ross. 2012. China's internet: Topology mapping and geolocating. In *2012 Proceedings IEEE INFOCOM*. IEEE, 2531–2535.

[35] An Wang, Wentao Chang, Songqing Chen, and Aziz Mohaisen. 2018. Delving into internet DDoS attacks by botnets: characterization and analysis. *IEEE/ACM Transactions on Networking* 26, 6 (2018), 2843–2855.

[36] Yong Wang, Daniel Burgener, Marcel Flores, Aleksandar Kuzmanovic, and Cheng Huang. 2011. Towards {Street-Level}{Client-Independent}{IP} Geolocation. In *8th USENIX Symposium on Networked Systems Design and Implementation (NSDI 11)*.

[37] Zhiyuan Wang, Fan Zhou, Wenxuan Zeng, Goce Trajcevski, Chunjing Xiao, Yong Wang, and Kai Chen. 2022. Connecting the hosts: Street-level IP geolocation with graph neural networks. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*. 4121–4131.

[38] Aviram Zilberman, Adi Offer, Bar Pincu, Yoni Glickshtein, Roi Kant, Oleg Brodt, Andikan Otung, Rami Puzis, Asaf Shabtai, and Yuval Elovici. 2024. A Survey on Geolocation on the Internet. *IEEE Communications Surveys & Tutorials* (2024). doi:10.1109/COMST.2024.3518398